RTC
The magazine of record for the embedded computing industry

Also in this issue

PC/104 Takes on High-End Apps

DSP for Control and Inspection

RTOS for DSP

March 2006 www.rtcmagazine.com

ETHERNET Adapts to the Needs of INDUSTRIAL CONTROL

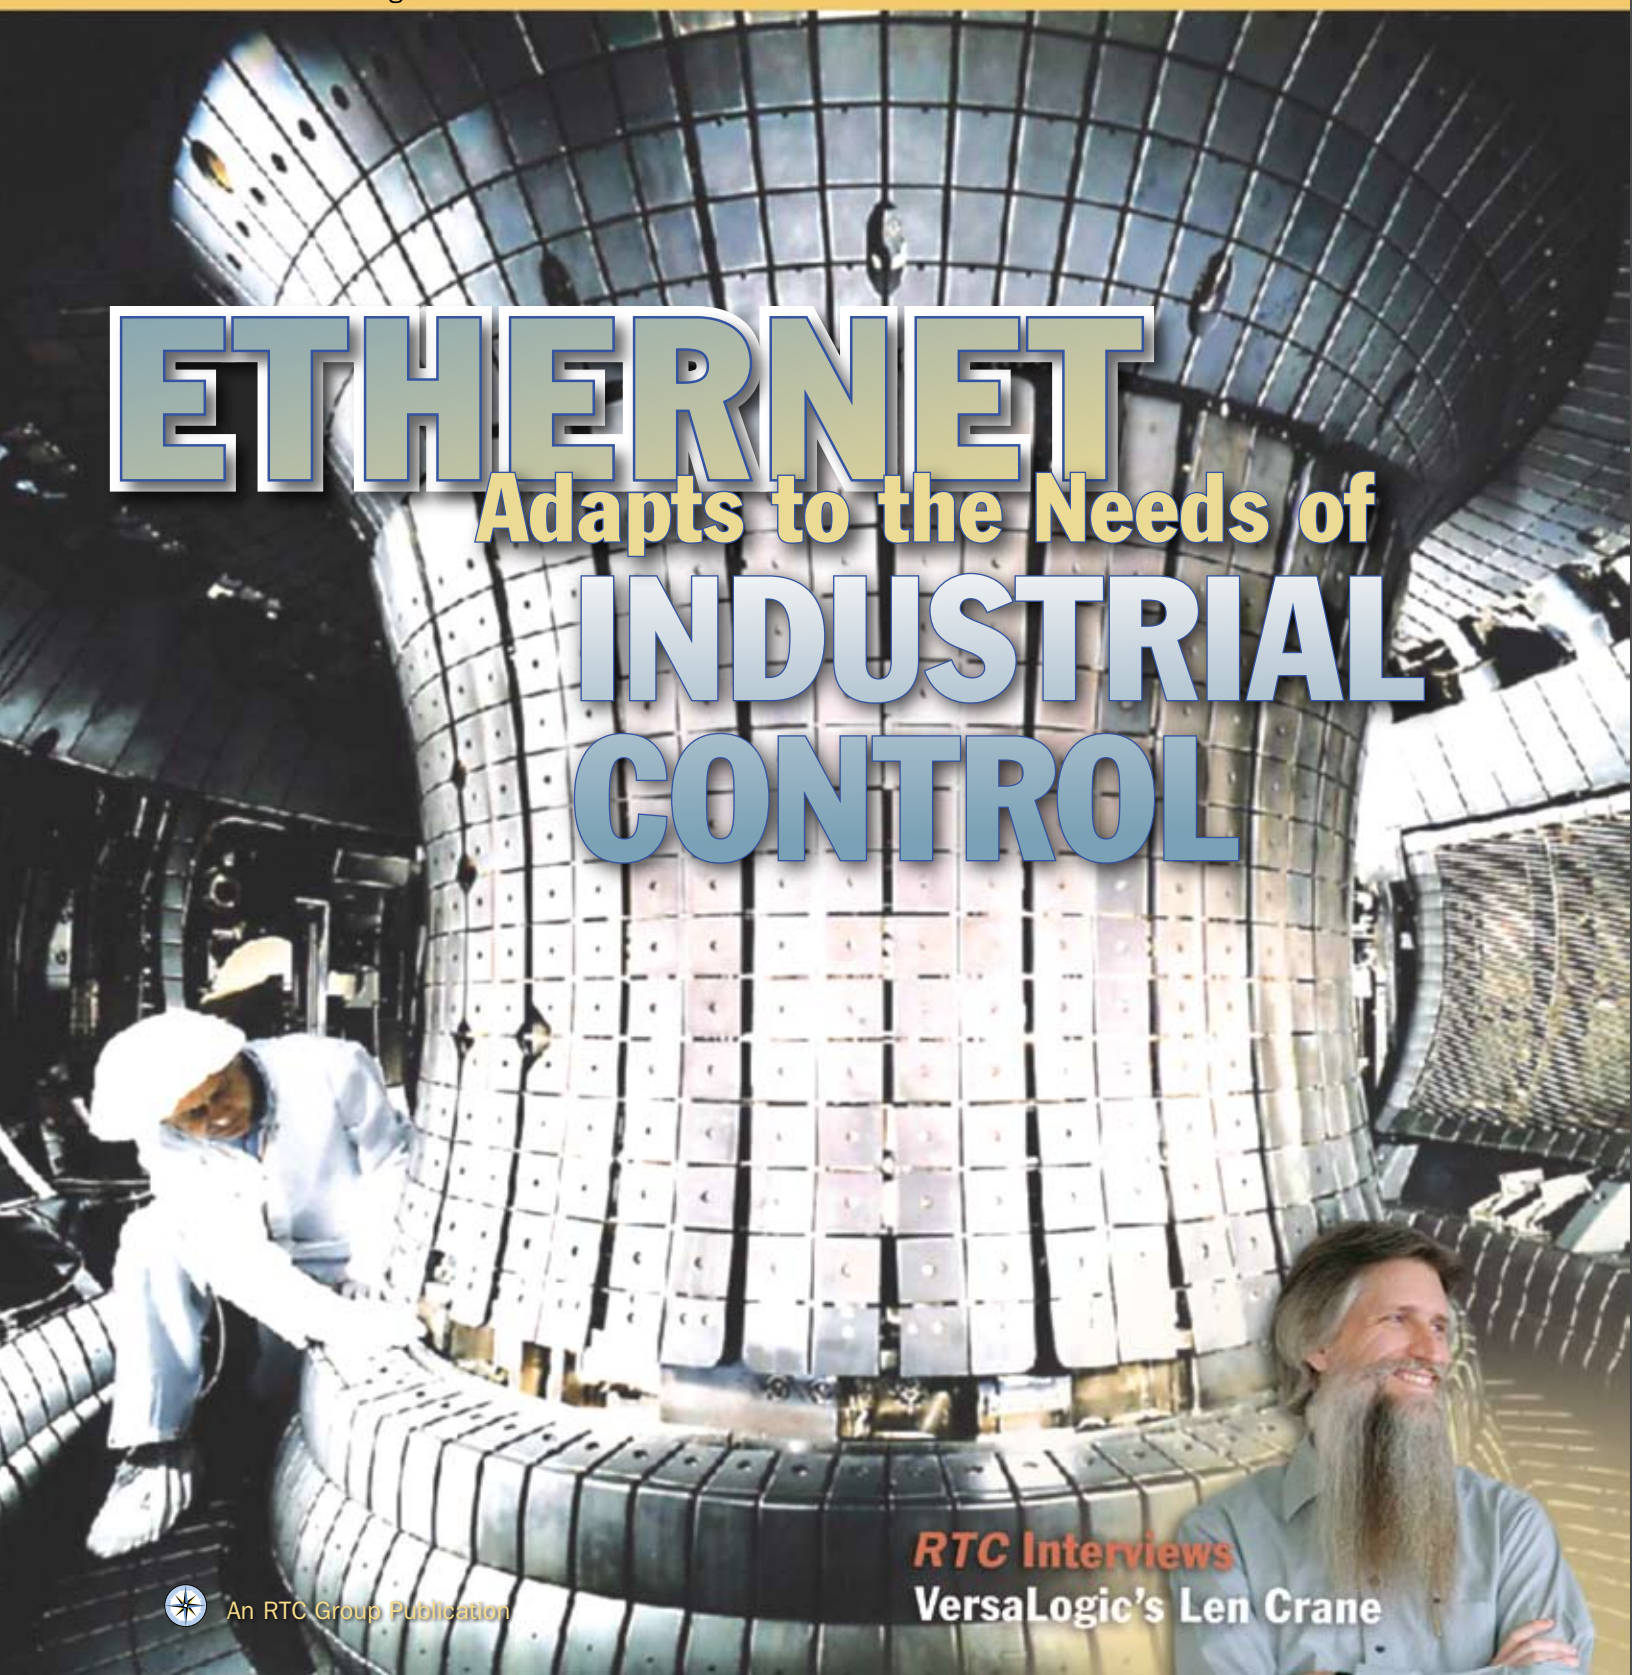

RTC Interviews
VersaLogic's Len Crane

An RTC Group Publication

Need reliability?
–40° to +85°

EPIC™ XE–900
1.0 GHz CPU

Features	XE–900	XE–800	XE–700
CPU	Via Eden	AMD Geode GXI	STPC
Clock speed	400 MHz; 733 MHz; 1.0 GHz	300 MHz	133 MHz
BIOS	General Software	Phoenix	Phoneix
DRAM support	to 256 MB	to 256 MB	32/64 MB
Compact/Flash	Type I or II	Type I or II	Type I or II
COM 1	RS–232	RS–232/422/485	RS–232
COM 2	RS–232	RS–232/422/485	RS–232/422/485
COM 3	RS–232	NA	RS–422/485
COM 4	RS–232	NA	RS–232
COM 5	RS–232/422/485	NA	NA
COM 6	RS–422/485/TTL	NA	NA
LPT1	0	0	1
EIDE	2	2	1
USB	2	6	2
CRT	1600 x 1200	1280 x 1024	1280 x 1024
Flat panel	LVDS	yes	yes
Digital I/O	24–bit prog.	48–bit prog.	24–bit prog.
Ethernet	10/100 Base–T	Dual 10/100 Base–T	10/100 Base–T
Expansion	PC/104 & *Plus*	PC/104 & *Plus*	PC/104
Power	3.6A operating	1.6A max.	1.6A max
Temp. range	–40° to 70/85° C	–40° to 80° C	–40° to 80/85° C
Shock/vibration	40/5g	40/5g	40/5g

Need Linux, QNX, Windows®?
Try our OS EMBEDDER™ KITS

Our kits are the shortest path to a successful OS on an Octagon embedded computer.

- Pick your Octagon SBC
- Pick the OS you prefer: Linux, Windows, QNX

Octagon delivers a high performance, total solution.

Typical Linux kit includes:

- *Target CPU card*
- *Preloaded OS image on 256 MB industrial CompactFlash*
- *256 MB SO–DIMM module*
- *Interface cables*
- *Hard copy of manual*
- *Mouse*
- *CPU OS bootable CD*
- *Optimized OS version*
- *Full driver support for on–board hardware*
- *X–Windows support*
- *Example applications and source code*
- *Extra documentation*

Need PC/104 expansion?
Try our XBLOKs®

XBLOKs offer the best compromise in cost and function for both PC/104 and PC/104-*Plus*. Only 44% the size of a standard PC/104 card, you can add two functions to your system but increase the stack height by only one level. –40° to 85° C. Heat diagram shows enhanced cooling.

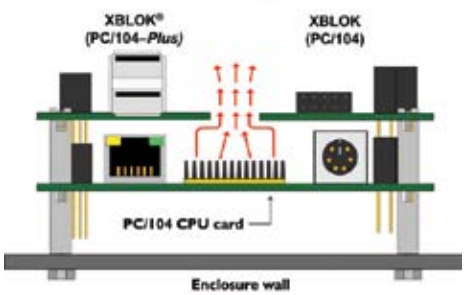

X–SRAM–2 MB
- 2 MB high speed, SRAM
- Read and write at full bus speed
- Pointers to memory saved if CPU resets or loses power

X–DIO–48 bit programmable digital I/O
- 48 digital I/O, 5V compatible
- Source and sink 16 mA per output
- Direct connection to opto-module racks

X–COM–2 dual UART
- Up to 230.4 kBaud data rate
- Supports RS–232/422/485
- RS–485 fault protected to ±60V

X–LAN–1 Ethernet LAN
- 10/100 Base–T, Intel 82551ER
- Fully plug–n–play
- High performance, PCI bus interface

X–USB–4 quad USB 2.0
- Speeds up to 480 mbps
- Mix and match USB 1.1 and 2.0
- Current–limited ports can supply 500 mA to external devices

Need a fanless system?
NEW CONDUCTION COOLING SYSTEM

Designed for the **XE-900**, our conduction cooling system eliminates a fan even at 1.0 GHz.

For a full listing of Octagon Systems products, visit us at
www.octagonsystems.com

OCTAGON SYSTEMS

Selecting the right Pentium® M-based computer has never been easier.

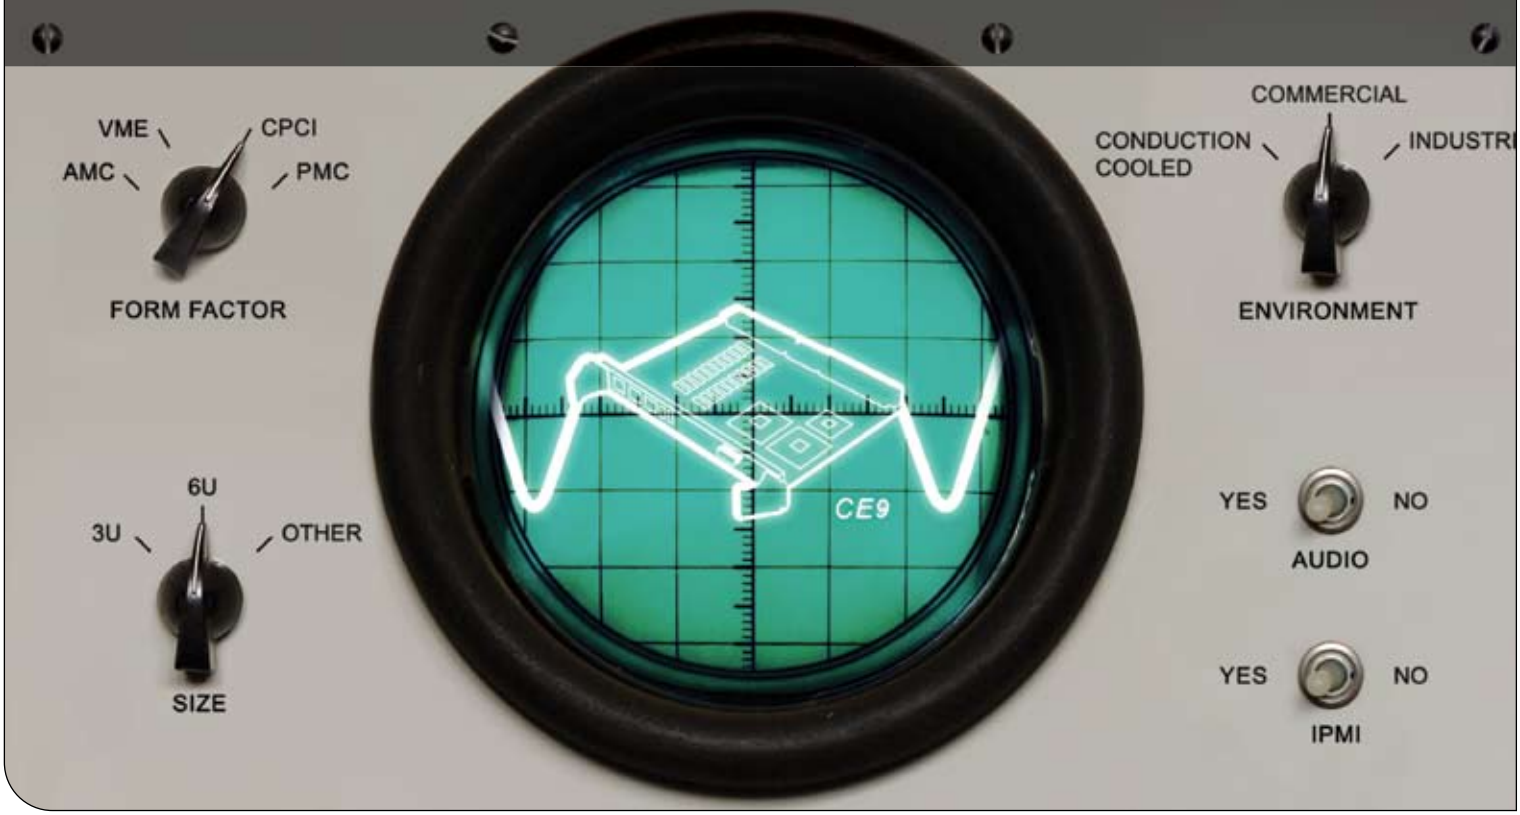

SBS knows Intel® Pentium® M processors. We have the CompactPCI®, VME, PMC or AdvancedMC™ single board computer you need.

Intel® Communications Alliance
Associate Member
SILVER

RoHS compliant 2002/95/EC

CE9 available with ROHS Compliance

SBS HAS NINE INTEL® PENTIUM® M processor-based single board computers for you to choose from. So it's likely that we have exactly the board you need. And we offer many of our computers in several configurations, with various I/O types and chipsets designed for specific applications, from industrial to defense to telecommunications.

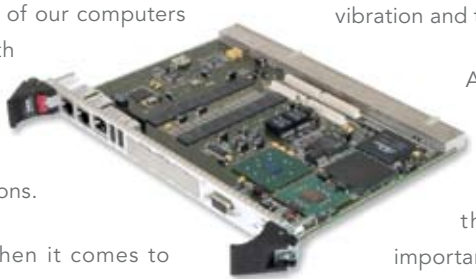

We also offer you options when it comes to processor speeds and types, as well as memory configurations. Depending on the individual product, SBS can often give you the choice of a standard-duty product, an industrial-duty product with extended temperature ranges, or a rugged-duty option designed to survive extreme shock, vibration and temperature variations.

After nearly two decades of designing embedded computing products, SBS can also offer you something that is perhaps even more important than product selection, and that is the benefit of our experience. So when it's time to select an Intel® Pentium® M processor-based single board computer, select SBS.

SBS Technologies

SBS knows. Find the compact Intel® Pentium® M processor board you're looking for at **www.sbs.com/pentiumm** or call **800.SBS.EMBEDDED**

Departments

9 **Editorial:** Telecom Is Not Dead—It Is Undead

10 **Industry Insider**

64 **Products&Technology**

Features

Technology in Context Industrial Ethernet

14 **Real-Time Ethernet is Reaching the Field for Industrial Control**
Frank Hansen, Altera

20 **EtherCAT: A Real-Time Industrial Fieldbus on Ethernet**
Paul Fischer, TenAsys

25 **DDS Over Ethernet Simplifies Industrial Automation**
Bob Kindel, Real-Time Innovations

Solutions Engineering PC/104 Tackles High-End Apps

31 **EPIC Express: A Bridge to the Future for High-Performance I/O**
Susan Wooley, Micro/sys

37 **PCI Express Stacks Up in High-End Embedded Applications**
Colin McCracken, Ampro Computers

Industry Insight DSP for Control and Inspection

44 **DSPs, Microprocessors and FPGAs in Control**
Shelley Gretlein, Gerardo Garcia and Joel Sumner, National Instruments

50 **DSPs + FPGAs Provide Flow Cytometry and Cell Sorting Solution**
Kristen K. Zaffini, BittWare

Executive Interview

55 ***RTC* Interviews Len Crane, President and CEO, VersaLogic**

Software & Development Tools RTOS for DSP

61 **Evaluating Real-Time Operating Systems for DSP**
Tom Barrett, Quadros Systems

Industry Watch

71 **PMCs: They're the Real Thing**
David Compston, Radstone Embedded Computing

Photo Credit: The containment vessel of the ASDEX Tokomak fusion reactor at the Max Planck Institute for Plasma Physics. The reactor uses an Ethernet Powerlink communications network to set up different experiments. Courtesy, Max Planck Institut für Plasmaphysik.

RTC

www.rtcmagazine.com

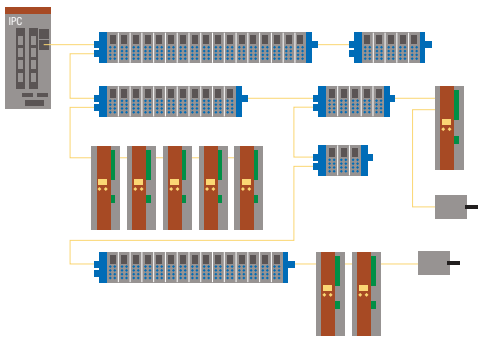

EtherCAT tree topology. • **Pg. 14**

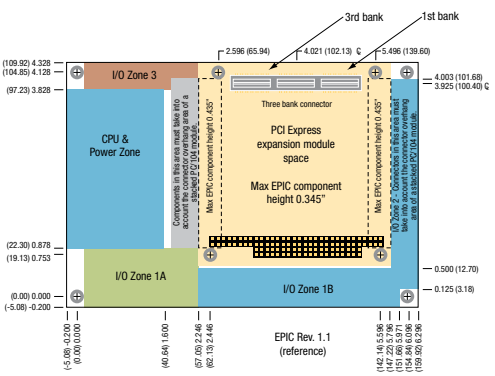

The EPIC Express form-factor defines a 1-bank and 3-bank PCI Express expansion connector. • **Pg. 37**

VITA 31 VME Board Targets Rugged Mission-Critical Apps • **Pg. 64**

When Life Is On The Line
Operate and survive under the most extreme conditions with ruggedized E-Disk® solid-state flash drives and network storage solutions. BiTMICRO's cutting-edge storage technologies offer utmost reliability, optimum data security and unmatched performance.
Ethernet | Fibre Channel | SCSI | IDE/ATA
USB | FireWire | cPCI VME | SATA | iSCSI
PCI-X | PCI Express | SAS | Infiniband
BiTMICRO
ULTIMATE STORAGE SOLUTIONS™
BiTMICRO Networks, Inc.
45550 Northport Loop E
Fremont, CA 94538-6481
www.bitmicro.com
info@bitmicro.com
510-743-3475

Rugged CPU Boards and SBCs
Rugged CPU Boards and SBCs in CompactPCI, 3,5", VME, EPIC, ATX, Mini ITX, PICMG, PC/104 and PC/104-Plus formats.
• Intel® Pentium® M up to 2.26 GHz
• Operating temperature range from –40°C to +85°C
• Reliable in harsh environment
• Extensive communication capabilities
• Low power consumption
Embedded Systems Conference
ESC Silicon Valley
San Jose, CA
Booth № 5001
info@fastwel.com • www.fastwel.com
Fastwel
Creating the Future!

Publisher

PRESIDENT
John Reardon, johnr@rtcgroup.com

VICE PRESIDENT, European Operations
Zoltan Hunor, zoltanh@rtcgroup.com

EDITORIAL DIRECTOR/ASSOCIATE PUBLISHER
Warren Andrews, warrena@rtcgroup.com

Editorial

EDITOR-IN-CHIEF
Tom Williams, tomw@rtcgroup.com

SENIOR EDITOR
Ann Thryft, annt@rtcgroup.com

MANAGING EDITOR
Marina Tringali, marinat@rtcgroup.com

STAFF EDITOR
Craig Choisser, craigc@rtcgroup.com

COPY EDITOR
Rochelle Cohn

Art/Production

CREATIVE DIRECTOR
Jenna Hazlett, jennah@rtcgroup.com

GRAPHIC DESIGNER
Melissa Gaeta, melissag@rtcgroup.com

PRODUCTION DESIGNER
Kirsten Wyatt, kirstenw@rtcgroup.com

DIRECTOR OF WEB DEVELOPMENT
Jan Garcia, jang@rtcgroup.com

WEB DEVELOPER
Marke Hallowell, markeh@rtcgroup.com

Advertising/Web Advertising

CALIFORNIA COASTAL ADVERTISING MANAGER
Diana Duke, dianad@rtcgroup.com
(949) 226-2011

WESTERN REGIONAL ADVERTISING MANAGER
Lea Ramirez, lear@rtcgroup.com
(949) 226-2026

EASTERN REGIONAL ADVERTISING MANAGER
Nancy Vanderslice, nancyv@rtcgroup.com
(978) 443-2402

EMEA SALES MANAGER
Anna Brown, annab@rtcgroup.com
(949) 226-2025

ASIAN ADVERTISING MANAGER
John Koon, johnk@rtcgroup.com
(858) 755-4999

Billing

Maggie McAuley, maggiem@rtcgroup.com
(949) 226-2024

To Contact *RTC* magazine:

HOME OFFICE
The RTC Group, 905 Calle Amanecer, Suite 250, San Clemente, CA 92673
Phone: (949) 226-2000 Fax: (949) 226-2050, www.rtcgroup.com

EASTERN SALES OFFICE
The RTC Group, 96 Dudley Road, Sudbury, MA 01776
Phone: (978) 443-2402 Fax: (978) 443-4844

Editorial Office
Warren Andrews, Editorial Director/Associate Publisher
39 Southport Cove, Bonita, FL 34134
Phone: (239) 992-4537 Fax: (239) 992-2396

Tom Williams, Editor-in-Chief
245-M Mt. Hermon Rd., PMB#F, Scotts Valley, CA 95066
Phone: (831) 335-1509 Fax: (408) 904-7214

Ann Thryft, Senior Editor
15520 Big Basin Way, Boulder Creek, CA 95006
Phone: (831) 338-8228

Published by The RTC Group
Copyright 2006, The RTC Group. Printed in the United States. All rights reserved. All related graphics are trademarks of The RTC Group. All other brand and product names are the property of their holders.

Whatever Your Backplane

or lack thereof...

DPM

- VME64, VITA 31.1
- Pentium M to 1.8 GHz
- 855 GME chipset with embedded graphics
- Low power consumption
- -40°/+71°C operation

CPM1

- CompactPCI, PICMG 2.16
- Pentium M to 1.8 GHz
- 855 GME chipset with embedded graphics
- Two PMC sites

KPM-S1

- Custom embedded
- Pentium M 2 GHz
- 7520 chipset with up to 8 GB DDR 400
- ATI M26 Mobility Radeon graphics on PCI Express

Whatever Your Environment

Conduction cooled versions of the VME and CompactPCI/PICMG 2.16 single board computers

- Extended temperature to -40°/+85°C
- Tested to MIL STD 810 and 901D
- Tested to high shock/vibration
- Conformal coating available
- Customized interfaces supprted

Let Dynatem provide your next embedded processor solution

DYNATEM

www.dynatem.com 800.543.3830 949.855.3235 sales@dynatem.com

Dynatem, Inc., 23263 Madero, Suite C, Mission Viejo, CA 92691

GE Fanuc Automation

Embedded Performance.

Looking for an embedded computing solution that gives you a tremendous advantage over your competition? Look no further than GE Fanuc Embedded Systems.

Featuring a comprehensive offering that includes Intel-based SBCs and complete I/O systems, industry-leading communications technology, rugged flat panel monitors and computers and more,

GE Fanuc Embedded Systems can support your full range of embedded computing needs to solve your greatest challenges. From standard product requests to a solution that is quickly and fully customized to your specific application, GE Fanuc Embedded Systems has the breadth, depth and support capabilities to provide a serious boost to your performance.

Learn more at www.gefanuc.com/embedded

Talon 8400
Self Contained
Embedded PC

- Intel 733 MHz Celeron M processor
- 256 MB of PC100/PC133 SDRAM
- Dual Ethernet ports
- Two Cardbus/PCMCIA slots
- Three PC/104-Plus slots
- 20 GB hard drive
- Integrated power supply
- Fanless operation
- Windows XP Pro OS
- Extended temperature and hazardous area models

CPCI-7808
Intel Pentium M CompactPCI
Single Board Computer

- PICMG 2.16/2.9 compliant
- Processor speeds up to 1.8 GHz
- Up to 2 GB DDR SDRAM
- Dual PMC sites
  - 64-bit/66 MHz site
  - 32-bit/33 MHz site
- Dual 10/100/1000 Ethernet Interface
- Dual 16550-compatible serial ports
- Three USB 2.0 ports
- Serial ATA interface
- Up to 1 GB CompactFlash
- OS support for Windows XP, Windows 2000, QNX, Linux, and VxWorks

PMC696
Intelligent Dual Gigabit Ethernet
PCI Mezzanine Card

- On-board high performance RISC processor
  - Up to 128 Mbytes of ECC DDR SDRAM
  - Dedicated TCP off-load engine
  - Built-in Ethernet failover capability
- 64 bit/66MHz PCI-X internal high speed bus
- Dual Gigabit Ethernet options
  - Fiber, 1000BaseFX
  - Copper, 10/100/1000BaseTX
- Trunking capabilities

CP920
CompactPCI Managed
Gigabit Ethernet Switch

- PICMG® 2.16 compliant
- Layer 2/3/4 switching
- Twenty-four 10/100/1000 Ethernet ports
- PICMG® 2.9 Rev 1.5 IPMI compliant
- PICMG® 2.1 Rev 2.0 hot swap compliant
- 802.1p, 802.1Q VLAN, deep packet filtering, link aggregation, Rapid Spanning Tree (802.1w, 802.1d), broadcast storm control, port mirroring
- Conduction cooled model available
  - Twelve 10/100/1000 Ethernet ports

GE FANUC

Embedded Systems

©2006 GE Fanuc Automation. All rights reserved.

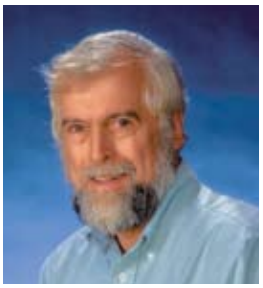

Telecom Is Not Dead—It Is Undead

by Tom Williams, Editor-in-Chief

Remember AT&T? It was one rather unfondly known as "Ma Bell," a huge regulated monopoly that ran the nation's telecommunications network. At the same time it produced technology, services and communications, it set policy, pricing, controlled access and created a rate structure that rivaled the Kabbalah for incomprehensibility. AT&T was broken up in 1984 by the Federal Courts and split into the mother company and the so-called "Baby Bells," the regional phone companies. The breakup of AT&T was greeted with a mixture of praise and sorrow, but since then (maybe or maybe not as a direct result) we have witnessed the burgeoning of the Internet, the spread of wireless services and the rise of many players large and small in the telecom arena.

Now it appears that we are seeing one crabbed hand moving spider-like over the edge of the coffin while the other pulls the stake from the heart so that the specter can rise once again to roam the world feasting on the blood of the living. Does this mean the eventual consolidation of the telecommunications market into one or two major players that, if not undead, will be unregulated? Will any combination of garlic or crucifixes be able to stop such a development?

Since the breakup of AT&T, quite a number of players have entered the market ushering in an era of competition, innovation and lower prices for consumers. SBC, the largest of the Baby Bells, then moved to acquire its former parent company late last year and has just recently adapted the AT&T name and logo. This sigil of the House of the Impaler was on my most recent phone bill. Now, this resuscitated AT&T has agreed to buy BellSouth in a $67 billion stock deal. BellSouth and AT&T were joint owners of Cingular Wireless, which will also become a wholly owned part of the new entity. As with the breakup, this development is being met with both optimism and trepidation.

AT&T becomes the dominant carrier in 22 states and controls the majority of long-distance, wireless and high-speed digital subscribers. The only serious U.S. competitor of any size is Verizon. In one sense, this appears to be a major blow in the ongoing battle between the telcos and the cable companies, which among other things are vying to be the supplier of fiber to the home. In that rivalry, only one can prevail. Or, the winner to the home might be wireless, but that could run up against spectrum limitations. Ten Gbit/s bandwidth is what will be needed to supply the "Triple Play" combination of voice, data and IPTV. Another consideration, of course, is that the telcos are feeling the hot breath of foreign competitors and possible acquisitions and size may be a defense against that as well.

Of paramount importance is the question of the impact that such consolidation may have on consumers and the industries that hope to sell into the telecommunications market. It would be unfortunate if an unregulated duopoly were to put the brakes on what appears to be an era of exciting change, innovation and low prices. There is the school of thought that the large companies will be able to bring the new technologies and infrastructure that the end users are waiting for. However, the track record of the large telcos in the past is less than inspiring. Ma Bell was slow, stodgy and very proprietary. Some of us even remember the days when you couldn't attach user-owned equipment to the phone system. The PR line out of AT&T and Verizon is, of course, that this development will bring about great things for the customer. We're not going to buy that one just yet.

The other question is whether companies selling to telecommunications equipment manufacturers or those manufacturers themselves will be negatively impacted by re-consolidation. Some say there is economic incentive for the telcos to adapt open standards like ATCA and AMC, the SAF interface specs and Linux. They may or may not do so. But even if they do, they will probably manufacture their equipment in-house or outsource it to contract manufacturers. This may be a reason that one major supplier of equipment appears to be withdrawing from the fight by shutting down and merging its networks division into a government and enterprise division.

Our segment of the industry is going to have to be fast on its feet to keep up with developments. If the intent was to sell into a market composed of competing smaller companies whose interest was in building their value-add on open standards platforms, that logic might have to be re-thought. A market composed of only a couple of very large companies is a very different animal. It is probably a safe bet that this latest move is not yet the last, and a good deal of the battle is being fought with lobbyists in chic restaurants in Washington DC—on the proven experience that it's easier to buy a Congressman than to innovate technology. It will all bear watching to see where it leads—both from the perspective of our careers and our status as consumers.

INDUSTRY INSIDER

March 2006

Tom:

I want to congratulate you and *RTC* for your January issue. It was a marvelous piece of fiction, punctuated by your telecom editorial, and the letters from "Slick" Dick Somes, and Joe "The Dreamer" Pavlat of PICMG. I rarely read fiction, but found this issue most entertaining and comical. If this was anything, it was an admission that telecom in nearly all forms is dead, and signaled the extreme danger in following any of these faulty specifications.

I hope you will cover the Ruggedized AMC effort when it comes into the V-56 committee as well as you have covered the paranoid and defensive responses from PICMG about their AMC and µTCA junk-products falling apart in a benign shock&vibe test held in Somes' garage using his 1977 Gremlin as the vibration source. AMC and µTCA are completely unusable in anything BUT a benign environment, and they fit perfectly in the air-conditioned central office of a financially defunct telecom service provider. Misery does love company, as the saying goes.

Again, my sincere applause to *RTC* for lowering their standards of reporting and whitewashing a very serious issue, and creating a very funny piece. I'm sure it made a lot of well-informed engineers and users who read *RTC* laugh until their morning coffee shot out of their noses...

With my best regards,
Ray Alderman
Executive Director, VITA

PICMG Replies

Dear Editor:

The shock, vibration and seismic tests performed on the Advanced Mezzanine Cards reported to you recently were done in accordance with standards accepted worldwide for carrier grade equipment. Shock and vibration testing was done in accordance with IEC 61587-1 Performance Level DL1. The Seismic tests conformed to IEC 61587-2, ANSI T1.329 and NEBS Level 3, Zone 4, which is an operational test.

These tests are rigorous and equipment that passes can take a severe beating. On 9/11/2001 Verizon had several key switches within a one-block radius of the World Trade Center. When the Twin Towers collapsed, many of these wire center buildings suffered extensive and severe damage. However, much of the NEBS-certified equipment in these buildings continued to switch the traffic that was so vital to the rescue and recovery efforts. There was actually a Lucent 5ESS(R) switch in one of the sub-basements of the Twin Towers, and it continued to switch calls from under the rubble pile until its batteries finally ran down around a day after the collapse.

Joe Pavlat
President, PCI Industrial Computer Manufacturers Group

Mercury Partners with Image Management Specialist Acuo

Mercury Computer Systems has announced an extended partnership with Acuo Technologies, a provider of medical image management and archiving solutions, at the European Congress of Radiology. As part of this cooperation, Acuo, in collaboration with U.S. service providers, will include Visage PACS, Mercury's image management solution, as a bundled offering in its application service provider (ASP) offering in the United States.

In cooperation with other providers, Acuo Technologies has created a completely hosted solution for healthcare IT, including image-information and other management systems. Under this partnership, hospitals, clinics and physicians are able to access the Web-based software from Visage PACS locally via Acuo's peer-aware architecture, or centrally via data center-hosted solutions. Radiologists and clinicians can access the PACS solution at any time via the Internet and retrieve medical imaging data for interpretation, as well as archive image data in the data center. Visage PACS offers comprehensive workflow optimization and supports standards such as IHE and HL 7.

In addition, hospitals opting for this outsourcing solution can reduce the total cost of ownership. "Users only have to pay a usage fee, per study or monthly, for example. Further investments for IT infrastructure and software licensing are not needed," emphasized Shannon Werb, CTO at Acuo Technologies. "Our provider partners make the necessary infrastructure available in the data processing center and undertake the entire hardware and software maintenance. Thus, the users' expenditure of time and costs is dramatically reduced. Because data are accessed via VPN, ultimate data protection is ensured." Ten imaging centers and clinics in the United States have already signed up for the ASP services.

Celoxica Joins TI's Third-Party DSP Program

Celoxica has joined the Texas Instruments Incorporated TMS320 third-party program and presented its DSP-FPGA design solutions during the 2006 TI Developer Conference (TIDC) in Dallas, Texas. Celoxica's electronic system level (ESL) technology provides a C-based design and synthesis flow for DSP-FPGA system development. It leverages API specifications for DSP-FPGA communication and allows creation of applications and IP to a single set of standard interfaces independent of board level detail. Celoxica's design

Get Connected with companies mentioned in this article.
www.rtcmagazine.com/getconnected

"The QNX Neutrino microkernel OS has the perfect DNA for multi-core processors"

Dan Dodge. QNX CEO & CTO.
Pioneer in distributed and multiprocessor computing.

Introducing the **QNX® Momentics® development suite Multi-Core Edition,** the industry's most comprehensive software platform for multi-core systems. Powered by the massively scalable QNX Neutrino® RTOS, this fully integrated solution supports AMP, SMP, and BMP, a groundbreaking technology that simplifies code migration and future-proofs your designs for quad-core and beyond. It's the latest innovation from QNX Software Systems, the undisputed leader in multiprocessing technology.

Maximum Choice for Multi-Core

Only QNX gives you the power to choose the best multiprocessing model for your multi-core design:

	SMP	BMP	AMP
Seamless resource sharing	✓	✓	—
Scalable beyond dual-core	✓	✓	Limited
Mixed OS environment	—	—	✓
Dedicated processor by function	—	✓	✓
Inter-core messaging	Fast (OS primitives)	Fast (OS primitives)	Slower (application)
Thread synchronization between cores	✓	✓	—
Dynamic load balancing	✓	✓	—
System-wide debug & optimization	✓	✓	—

QNX Unlocks the Power of Multi-Core

Maximize performance. Eliminate complexity. Accelerate migration. Only QNX offers:

- **Asymmetric Multiprocessing (AMP)** for full developer control and fault tolerance
- **Symmetric Multiprocessing (SMP)** for maximum concurrency and scalability
- **Bound Multiprocessing (BMP)** for the fastest code migration and minimum design complexity
- **Transparent Inter-Processor Communication (TIPC)** protocol for seamless Linux connectivity
- **System tracing tools** for fast debugging and optimization of multi-core applications
- **Off-the-shelf BSPs** for multi-core platforms based on MIPS®, PowerPC®, and x86 architectures

Discover how Dan and the QNX team deliver the shortest migration path to multi-core. Call 1 800 676 0566 or visit **www.qnx.com/innovate.**

QNX SOFTWARE SYSTEMS

QNX, Momentics, and Neutrino are trademarks or registered trademarks of QNX Software Systems GmbH & Co. KG and are used under license by QNX Software Systems International Corporation. All other trademarks belong to their respective owners. 301786 MC339.05

solutions support other popular design environments such as TI's Code Composer Studio Integrated Development Environment, Mathworks' Simulink and Diamond from 3L.

"The C-based design flow for software and the programmable hardware components of a system broadens the reach of DSP-based solutions containing FPGA technology. Celoxica's design solutions enhance the development experience of our customers who need to reduce design time and accelerate complex algorithms," said Jennifer Skinner-Gray, worldwide DSP partner network manager, TI.

TI's TMS320 third-party program is an extensive collection of global DSP development support. With more than 650 independent companies and consultants, TI's customers have easy access to a broad range of application software, development hardware and software and consulting services.

GE Fanuc to Purchase Assets of Condor Engineering

GE Fanuc Embedded Systems announced that it has signed an agreement to purchase the assets of Condor Engineering, including test and simulation products, embedded boards and core technology to meet the data communications needs of commercial and military aircraft manufacturers and suppliers. According to Maryrose Sylvester, President and CEO of GE Fanuc Automation, "The innovations that Condor Engineering has brought to the burgeoning avionics industry are well recognized as quality solutions, and we plan to build on that reputation to further serve the industry."

Condor Engineering, located in Santa Barbara, California, brings a broad and international customer base within both the aerospace and defense industries. The company's product set includes powerful hardware interface solutions for application engineers, extensive high-level API libraries for software developers, GUI monitor and control programs for integrators and test engineers, and core logic for embedded designers.

This technology supports F-16, F-18, F-22, C-17, C-130, Airbus and Boeing aircraft. Test and Simulation products include COTS avionics communications solutions including CompactPCI, PCI, PMC and VME platforms. The transaction is expected to close in the near future, subject to customary closing conditions.

PLX Technology Design Contest for PCI Express

PLX Technology has announced the PLX PCI Express Design Contest, which gives engineers worldwide the opportunity to display their PCIe expertise and a chance to win a $1,500 prize. The contest calls for each participating designer to write an application note on a unique design using the PLX ExpressLane PEX 8111 PCIe-to-PCI bridging device, then submit it to PLX by May 26. All entries will be evaluated and judged by a panel of PLX engineers and managers. The winning design, which will be announced June 12, will garner its developer a $1,500 American Express Gift Card.

Full documentation on the PEX 8111 bridge—including product brief, data book, press announcements and design notes—can be found at www.plxtech.com/products/pci_express/PEX8111/default.asp.

Designers interested in taking part in the PLX PCI Express Design Contest simply need to:

1. fill out a completed entry form at www.plxtech.com/Design_Contest/design_contest_entryform.htm;
2. write an application note featuring the PEX 8111 device; and
3. submit the entry form and final application note to designcontest@plxtech.com by May 26, 2006 at 12:00 a.m. (PDT).

Full details of the contest, along with rules and regulations, are at www.plxtech.com/Design_Contest/design_contest_rules.htm. All participants will be notified of the judging panel's selection by June 12, 2006.

SOMA Brings Triple Play to the "Last Mile" with TI DSPs

Texas Instruments has announced that SOMA Networks is using TI's TMS320C5509A digital signal processors (DSPs), TMS320C5416 DSPs and TMS320C6415 DSPs within SOMAport 300 Wireless Broadband Gateway modems and SOMA Macro Base Stations, to bring its last-mile technology to customers around the world. TI supplies digital processors to meet the compact yet demanding component needs of SOMA Networks, from the ultra-low power C5509A DSP to the high-performance C6415 DSP.

SOMA Networks' SoftAir wireless systems have been enabling service providers around the world to implement standards-based solutions that meet their business requirements. The SoftAir platform delivers primary-line voice, multi-megabit broadband access and future 802.11 wireless access capabilities, to residential and small office subscribers. The 1 GHz TI C6415 DSP is included in the SOMA Macro Base Station, which serves a typical cell radius of up to ten miles and thousands of subscribers at 12 Mbits/s. Along with the C5509A DSP-enabled SOMAport 300 wireless gateway modem, service providers can offer fully converged voice and broadband services to a wide range of geographic regions.

In areas where installing copper or fiber is impractical, SOMA Networks offers service providers a complete turnkey wireless solution that can be up and running within a matter of weeks rather than years. The SOMA solution does not require line-of-sight access to a base station or an external antenna, greatly simplifying installation and deployment.

Fiber-to-the-Home Subscribers Increase 70% in the Last Third of 2005

Two leading telecommunications industry organizations, the Fiber-to-the-Home (FTTH) Council and the Telecommunications Industry Association (TIA), have announced new estimates showing FTTH deployments now passing over 3.6 million U.S. homes, an increase of almost 1 million homes in four months. The new figures, released prior to the FTTH Council's quarterly meeting in Washington, DC next week, show the most rapid growth of FTTH deployment to date, with more than 230,000 additional homes passed every month.

"Typically telecom construction slows in winter months," said Michael Render of market research firm Render, Vanderslice and Associates, which prepared the estimates. "To see acceleration at this time of year indicates increased preference for FTTH as an infrastructure of choice." The figures also show acceleration among homes connected by FTTH, which normally lags homes passed by several months. Over the last four months, the estimate shows homes connected with FTTH have increased by 70% to 548,000 subscribers from 322,000.

"This is a good sign, but we have to keep it up," TIA President Matt Flanigan said. "The nations that build extremely fast networks will gain very significant advantages. The United States must win the bandwidth race." Both organizations have urged policymakers to reduce barriers to next-generation broadband deployment, such as streamlining cable television franchising processes and allowing municipalities to build broadband networks.

Performance Beyond Limits

... ready for Extreme Embedded Computing?

ETXexpress™

INSERT-READY

ACCELERATED DESIGN

SCALABILITY

LONGEVITY

COST-EFFECTIVE

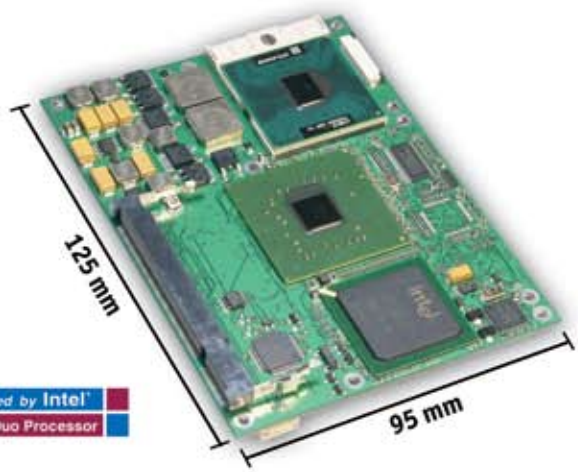

powered by Intel® Core® Duo Processor

ETXexpress-CD

- Highest performance state-of-the-art embedded module
- Intel® Core® Duo processor and Intel® 945GM chipset

ETXexpress products are next generation embedded modules based on the PICMG COM Express standard. ETXexpress provides the hightest performance and I/O bandwidth available in COMs.

- **PCI Express - the elemental data path**
- **Gigabit Ethernet - for high connectivity**
- **USB 2.0 - for fast periphery**
- **Serial ATA - for fast drives**
- **ACPI - for optimized power management**

Get ready. Get ETXexpress

Visit www.kontron.com/ETXexpress

ETXexpress www.etx-express.com

COM Express

Intel® Communications Alliance Premier Member

Computer-On-Modules | Blades & Mezzanines | CPU Boards | Systems | Mobile Rugged | Custom Solutions

1-888-294-4558 - sales@us.kontron.com — EMEA: +49 0 800 7253756 - ASIA: +886 2 2910 3532

kontron

Kontron and the Kontron logo are registered trademarks of Kontron AG. All other trademarks are the property of their respective owners. ©2005 Kontron America, Inc.

Real-Time Ethernet is Reaching the Field for Industrial Control

There are a number of solutions for implementing an industrial fieldbus using proven and widespread Ethernet technology. To do so, they must address real-time issues and cost. A number of cost-effective implementations can be achieved with FPGA technology.

by Frank Hansen
Altera

In 2003, about 30 years after the first developments on Ethernet by Xerox (later standardized as IEEE 802.3), the IEC in Geneva developed the standardization process for Real-Time Ethernet (RTE) to use Ethernet as a fieldbus alternative. Currently, industrial Ethernet is vigorously establishing itself in automation technology, and several promising approaches for RTE are forcing their way into the market. Besides the numerous devices that are already being employed in the industry, there are flexible and cost-effective FPGA-based implementations available. These will foster the further establishment of the standards while maintaining the highest possible levels of flexibility.

The RTE standardization process was driven by the development of industrial communication systems from different vendors. In 2001, the Austrian company B&R introduced the first approach for RTE with "Ethernet Powerlink." There are currently several organizations and many companies working to define an application protocol for Real-Time Industrial Ethernet. The Website http://www.real-time-ethernet.de provides a good overview on the different solutions (currently 17), vendors and supporting organizations (Table 1).

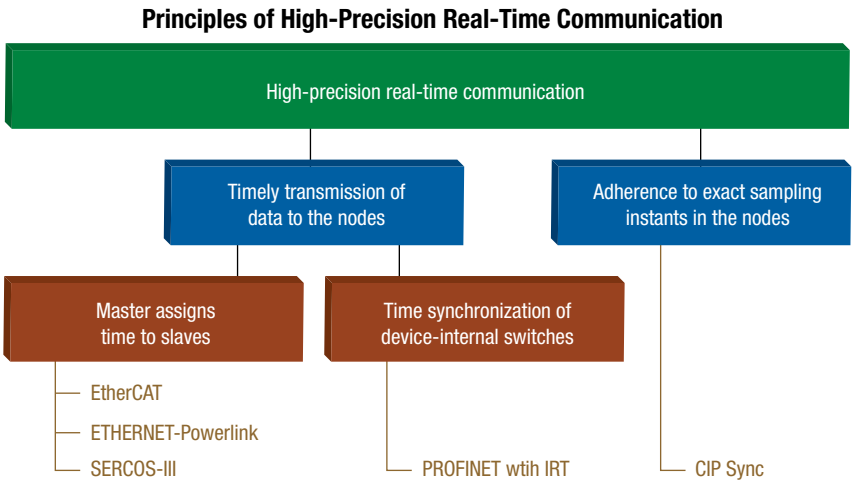

Figure 1 RTE demands for a deterministic timing approach (source: Prof. J. Schwager, University of Reutlingen, Germany).

What does "real time" mean in the context of the automation industry? It is specified as a defined or deterministic time in which a system or component can react to an external event. For communication tasks, not only the defined latency (cycle time) is important, but the jitter also has to be limited. Some automation or drive applications demand cycle times of less than 100 μs and jitter less than 1 μs. With standard Ethernet jitter, specifications of only 100 μs to 3 ms are possible. The various RTE approaches, such as Ethernet Powerlink, EtherCAT, Sercos III, ProfiNet and

Get Connected
with companies mentioned in this article.
www.rtcmagazine.com/getconnected

www.apacer.com

Expand Your Application Possibilities

With

+85°C

Apacer

-40°C

RoHS Compliant
Embedded Flash Solutions

See us at ESC Booth 3333

Extreme Reliability, Flexibility and Performance

Apacer has designed and engineered Industrial Grade Embedded Flash Memory products to provide the uncompromised reliability, flexibility and performance for applications in the severe extreme temperature and high vibration environments.

ADM II
ATA-Disk Module II

- IDE Interfaced
- In 40 & 44 Pins

AFD II
ATA Flash Drive II

- 1.8" & 2.5 HDD Format
- With High Capacity

Industrial CF II
Industrial CompactFlash Card II

- Industrial Grade CF
- With Hardware & Firmware Control

Global Presence

Apacer Memory America, Inc.
Tel: 408-232-1291
Email: embedded@apacerus.com
Website: www.apacer.com

Apacer
Access the best

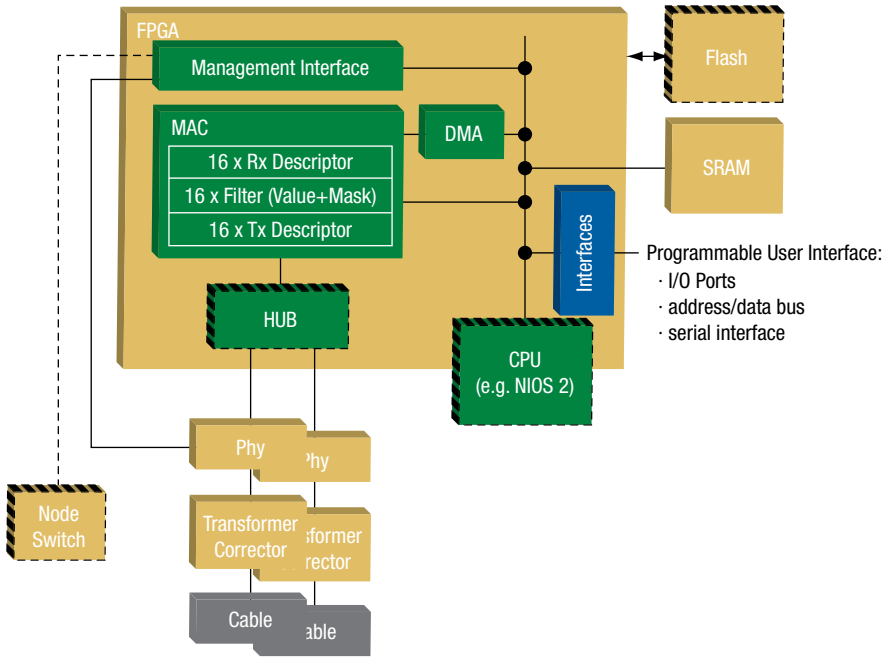

Figure 2 Block diagram of the EPSG-FPGA implementation with a Cyclone FPGA and, for example, a Nios II processor.

Ethernet/IP, differ significantly and are not compatible with one another.

The basic approach to ensure real-time communications is the use of a master to control the timing behavior of the data transfer to the slaves on the bus. Ethernet-Powerlink, EtherCAT and SERCOS-III are based on this principle. ProfiNet Isochrone Real Time (IRT) works with internal switches in all devices to handle the synchronization. Other approaches, like JetSync and Ethernet/IP CIP Sync, use time stamps for the synchronization.

The RTE systems have set the goal of enabling high-speed, real-time communication on the basis of the Ethernet transmission technique. To do this, all systems use specific real-time protocols and, for real-time communication, go without the TCP/IP protocol used in standard Ethernet (Figure 1). The following basic technologies and features of Ethernet Powerlink, EtherCAT, SERCOS III and, very briefly, Ethernet/IP and ProfiNet IO, are described. All are supported by industrial organizations and leading vendors.

Ethernet Powerlink

Ethernet Powerlink is an open real-time solution that was standardized by the Ethernet Powerlink Standardization Group (EPSG). Powerlink is compatible with the IEEE 802.3u standard for Fast Ethernet. In this way, it is possible to transmit all IP-based protocols, such as TCP, UDP and http, simultaneously with the hard real-time data. Ethernet Powerlink and EtherCAT were developed from the beginning as RTE solutions. So right from the start, the focus in the development of Ethernet Powerlink was on achieving very short cycle times of several hundred microseconds with very tight synchronization tolerances below 1 ms. These framework conditions ensure that Ethernet Powerlink can also be used to connect and synchronize very dynamic drive systems. Simultaneously, bandwidth is reserved to ensure a spontaneous exchange of data between the automation devices, using TCP/IP for instance. Up to 256 stations in the network can be handled, with jitter at all stations less than 1 μs.

Ethernet Powerlink uses isochronous Slot Communication Network Management (SCNM). With Ethernet Powerlink having the same topologies as Fast Ethernet (star, tree and line structures), up to 100m segment length and the usage of Category 5 patch cables is possible, both enabling cost-effective solutions compared to traditional fieldbus systems.

EtherCAT

EtherCAT is an open real-time solution that was developed by Beckhoff and is supported by the EtherCAT Technology Group (ETG). With EtherCAT, it is possible, for example, to transmit 1,000 distributed I/O data points in 30 μs or to synchronize 100 axes in 100 μs. EtherCAT is flexible with regard to the installation, supports line as well as tree or star topologies and does all this without active network components such as switches or hubs.

With EtherCAT, the costly Ethernet star topology can be replaced with a simple line or tree structure—no expensive infrastructure components are required. All types of Ethernet devices can be integrated via a switch or switch port. Where other real-time Ethernet approaches require special master or scanner cards, EtherCAT manages with very cost-effective standard Ethernet network interface cards (NICs).

RTE Solution	Organization	Special Ethernet Controller	Principle
Ethernet Powerlink	EPSG	No	Master/Slaves
EtherCAT	ETG	Yes	Master/Slaves
SERCOS-III	IGS	Yes	Master/Slaves
Ethernet/IP with CIPSync	ODVA	Yes	Time Stamps
Profinet/IRT	PNO	Yes	Internal Switches
Other RTE solutions: RTnet, JetSynC, PowerDNA, SynqNet, SynUTC, Precision Time Protocol (IEEE 1588), RTPS, Vnet/IP, EPA, SafetyNETp, Switch with Time Server, TCnet			

Table 1 Overview of the different available RTE solutions.

On the hardware side, EtherCAT technology is located in EtherCAT terminals. In contrast to bus terminals, where the fieldbus protocol data is converted within the bus coupler to the internal, fieldbus-independent terminal bus, the EtherCAT protocol remains fully intact down to the individual terminal. In addition to EtherCAT terminals with E-bus connection, the proven standard bus terminals with K-bus connection can also be connected via an EtherCAT bus coupler. This ensures compatibility and continuity with the prevalent system. Servo drives, fieldbus boxes, embedded PCs, Ethernet PCI cards and control panels with control functionality already offer inherent EtherCAT capability.

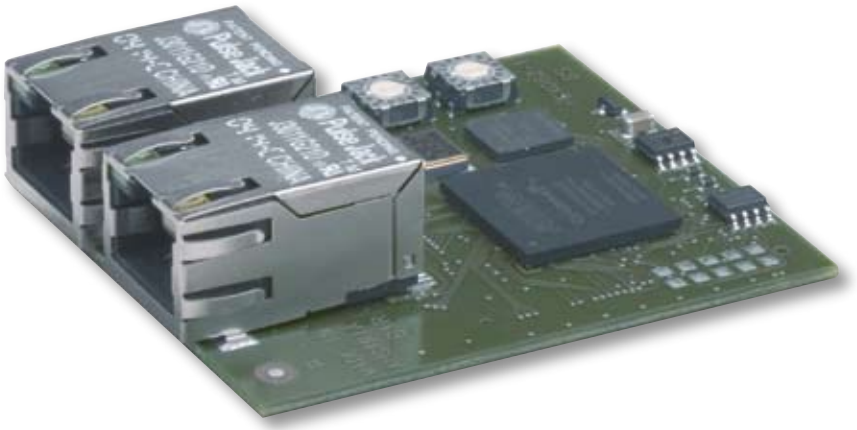

Figure 3 FPGA-based EPSG reference design (EX082).

SERCOS

The SERCOS Interface Interest Group (IGS) is a global association that is responsible for the commercialization, certification, further technical development and standardization of the SERCOS interface real-time communication standard.

Since its inception in the 1980s, the SERCOS interface has built support as a standardized (IEC61491) means of communication between motion controllers and digital drives, I/O and sensors. A number of control and drive manufacturers worldwide offer SERCOS interface products.

The challenge was to create the best of both worlds and combine the high bandwidth and low cost of the non-deterministic Ethernet with the deterministic performance and standardized parameters of the SERCOS interface. The SERCOS-III solution places the standard Ethernet TCP/IP under control of the motion bus and uses lower-cost Ethernet hardware and twisted-pair copper cable. This maintains the deterministic motion control of the SERCOS interface, allows links to the existing manufacturing communications infrastructure, provides for the possibility of new features and lowers hardware costs.

SERCOS-III offers the added benefits of providing integration of IP protocols, new cross communication between slaves, synchronization of several motion controls, line or ring topology and fault tolerance in case of a break in the ring. A low-cost FPGA communication controller can be used to handle all sequencing and synchronization tasks, relieving the host microprocessor of these time-intensive functions. With the SERCOS-III IP core, manufacturers of components and systems can combine the SERCOS-III hardware and their own logic components in one common FPGA. This approach lowers cost, maintains determinism and provides transmission speed to 100 Mbits/s.

For the sake of completeness, the main and earliest RTE approaches are Ethernet/IP and ProfiNet. The Industrial Ethernet Protocol (Ethernet/IP) was developed by ODVA and Rockwell Automation and is based on a protocol extension on Layer 7. For the application protocol, the Common Industrial Protocol (CIP), which is known as the application layer from DeviceNet and ControlNet, is used. Ethernet/IP was extended with CIP Sync to provide real-time data transfer. All RTE components in the network must have implemented a synchronization device based on the IEEE 1588 standard.

ProfiNet was developed by Siemens and the ProfiBus association (PNO). ProfiNet offers three performance classes: standard communication (using TCP and UDP), real time (RT) and IRT. To handle "hard" real-time applications ProfiNet offers IRT integrated in special ASICs. These ASIC solutions integrate not only the Ethernet controller but also a switch to support line infrastructures.

Example Implementation: Ethernet Powerlink

Unlike the other RTE approaches, Ethernet Powerlink can be implemented with any standard Ethernet chip and does not need a special Ethernet controller. Approaches developed in-house on an ASIC basis could ultimately be somewhat faster; nevertheless, as many proprietary fieldbus developments have shown, limitation to a single chip manufacturer, restrictive licensing or patent issues, as well as availability issues, can turn into a disadvantage. For this reason, one of Ethernet Powerlink's most important advantages is the fact that the protocol can be implemented on any standard Ethernet hardware and on any corresponding chip.

The Ethernet Powerlink devices already being used today are based on different chip architectures that range from designs with standard processors to programmable FPGA hardware. The device manufacturer can decide which chip best meets its requirements. A solution with an optimized, license-free Powerlink interface, hub logic and a processor core can be realized, for example, in a cost-optimized FPGA plus one external PHY. In this way, manufacturers of cost-sensitive devices (like small sensors) have access to a universal, flexible architecture that can be adapted to the needs of the application. Altera and other Ethernet Powerlink partners offer a reference design for plug-and-play solutions. Very few device manufacturers have the time and the in-house resources to acquire the expertise for setting up and fully implementing a communication protocol such as Powerlink. It is usually better to concentrate on the primary competencies in the device function and ensure a high-speed, high-quality network connection. Several companies offer custom services and components that make

Powerlink fast and easy to implement. The customer can take advantage of turnkey software stacks and reference boards for different chip architectures.

ASIC, Microcontroller or FPGA?

In principle, there are different possibilities for implementing Ethernet Powerlink. It can be completely implemented in software without limiting the real-time capabilities. In a typical new implementation in an embedded design, both the Ethernet Powerlink communication and the corresponding applications usually run on the microcontroller. A system that already exists can also be easily equipped with Ethernet Powerlink functionality by adding a special microcontroller that takes over the Powerlink communication. This additional microcontroller can also be in the form of an FPGA with a microcontroller core and the corresponding Ethernet controller. In this case, the Powerlink protocol software is executed completely by the additional microcontroller or FPGA so that the application on the system CPU is not affected (Figure 2).

Solutions that are completely based on an FPGA are particularly flexible. This solution particularly benefits manufacturers of actuators and sensors. These companies can then react quickly to changing market requirements while significantly reducing front-end costs. This provides suppliers of cost-critical automation components with a universal, flexible architecture that they can adapt to the given application. The FPGA-based solution offers several advantages for the establishment of the EPSG standard:

- Universal EPSG interface
- Reprogrammability
- Optimized frame filter
- Flexible processor core for applications
- Configurable peripheral functions and interfaces
- Reduced costs

The first FPGA-based Ethernet Powerlink implementation (EX082) was developed by the EPSG. A single Cyclone FPGA contains a license-free, optimized Ethernet Powerlink interface, the hub logic and a processor core. This makes a low-priced and future-proof solution available for realizing EPSG-capable Ethernet nodes (Figure 3). The board can be configured with a boot PROM or via the serial interface. Sixteen optimized hardware filters ensure trouble-free sending and receiving operations with the powerful 16-channel MAC unit. A 32-bit soft-core processor can be used for processing the Powerlink protocol and the application. Besides the two RJ45 jacks and the two PHY components, there are also the node switches and a LED status display on the board. Interfaces to the address/data bus, external flash memory and serial I/Os that supplement the functional range of the reference design beyond the EPSG evaluation boards are available.

Altera
San Jose, CA.
(408) 544-7000.
[www.altera.com].

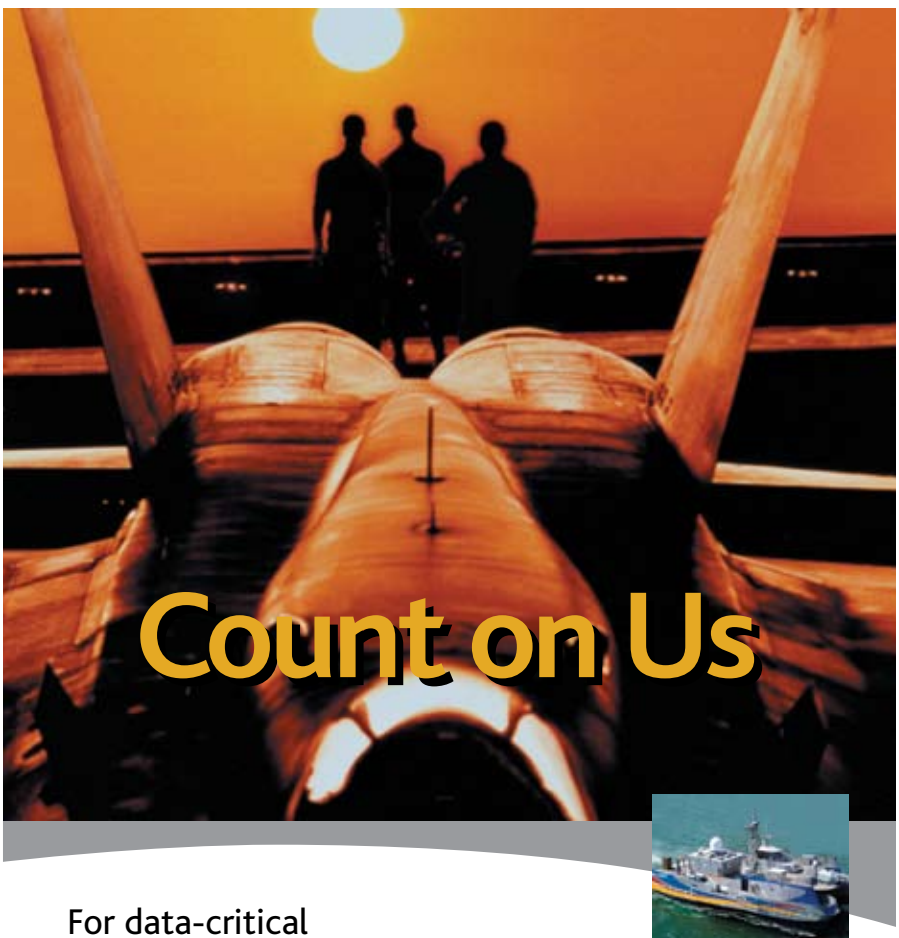

NDDS powers Lockheed Martin's Sea SLICE

For data-critical networked applications

At the end of the day, mission success relies on the ability to respond to changes in an instant. Designers of mission-critical combat systems count on Real-Time Innovations (RTI) for fast and reliable standards-based software to communicate real-time data over a network. RTI's NDDS middleware supports the DDS standard and is widely used in military and aerospace applications today.

Interested in learning how NDDS can work in your application? Download *Using the DDS standard for High-Reliability Applications*, free at www.rti.com/rtc06

RTI

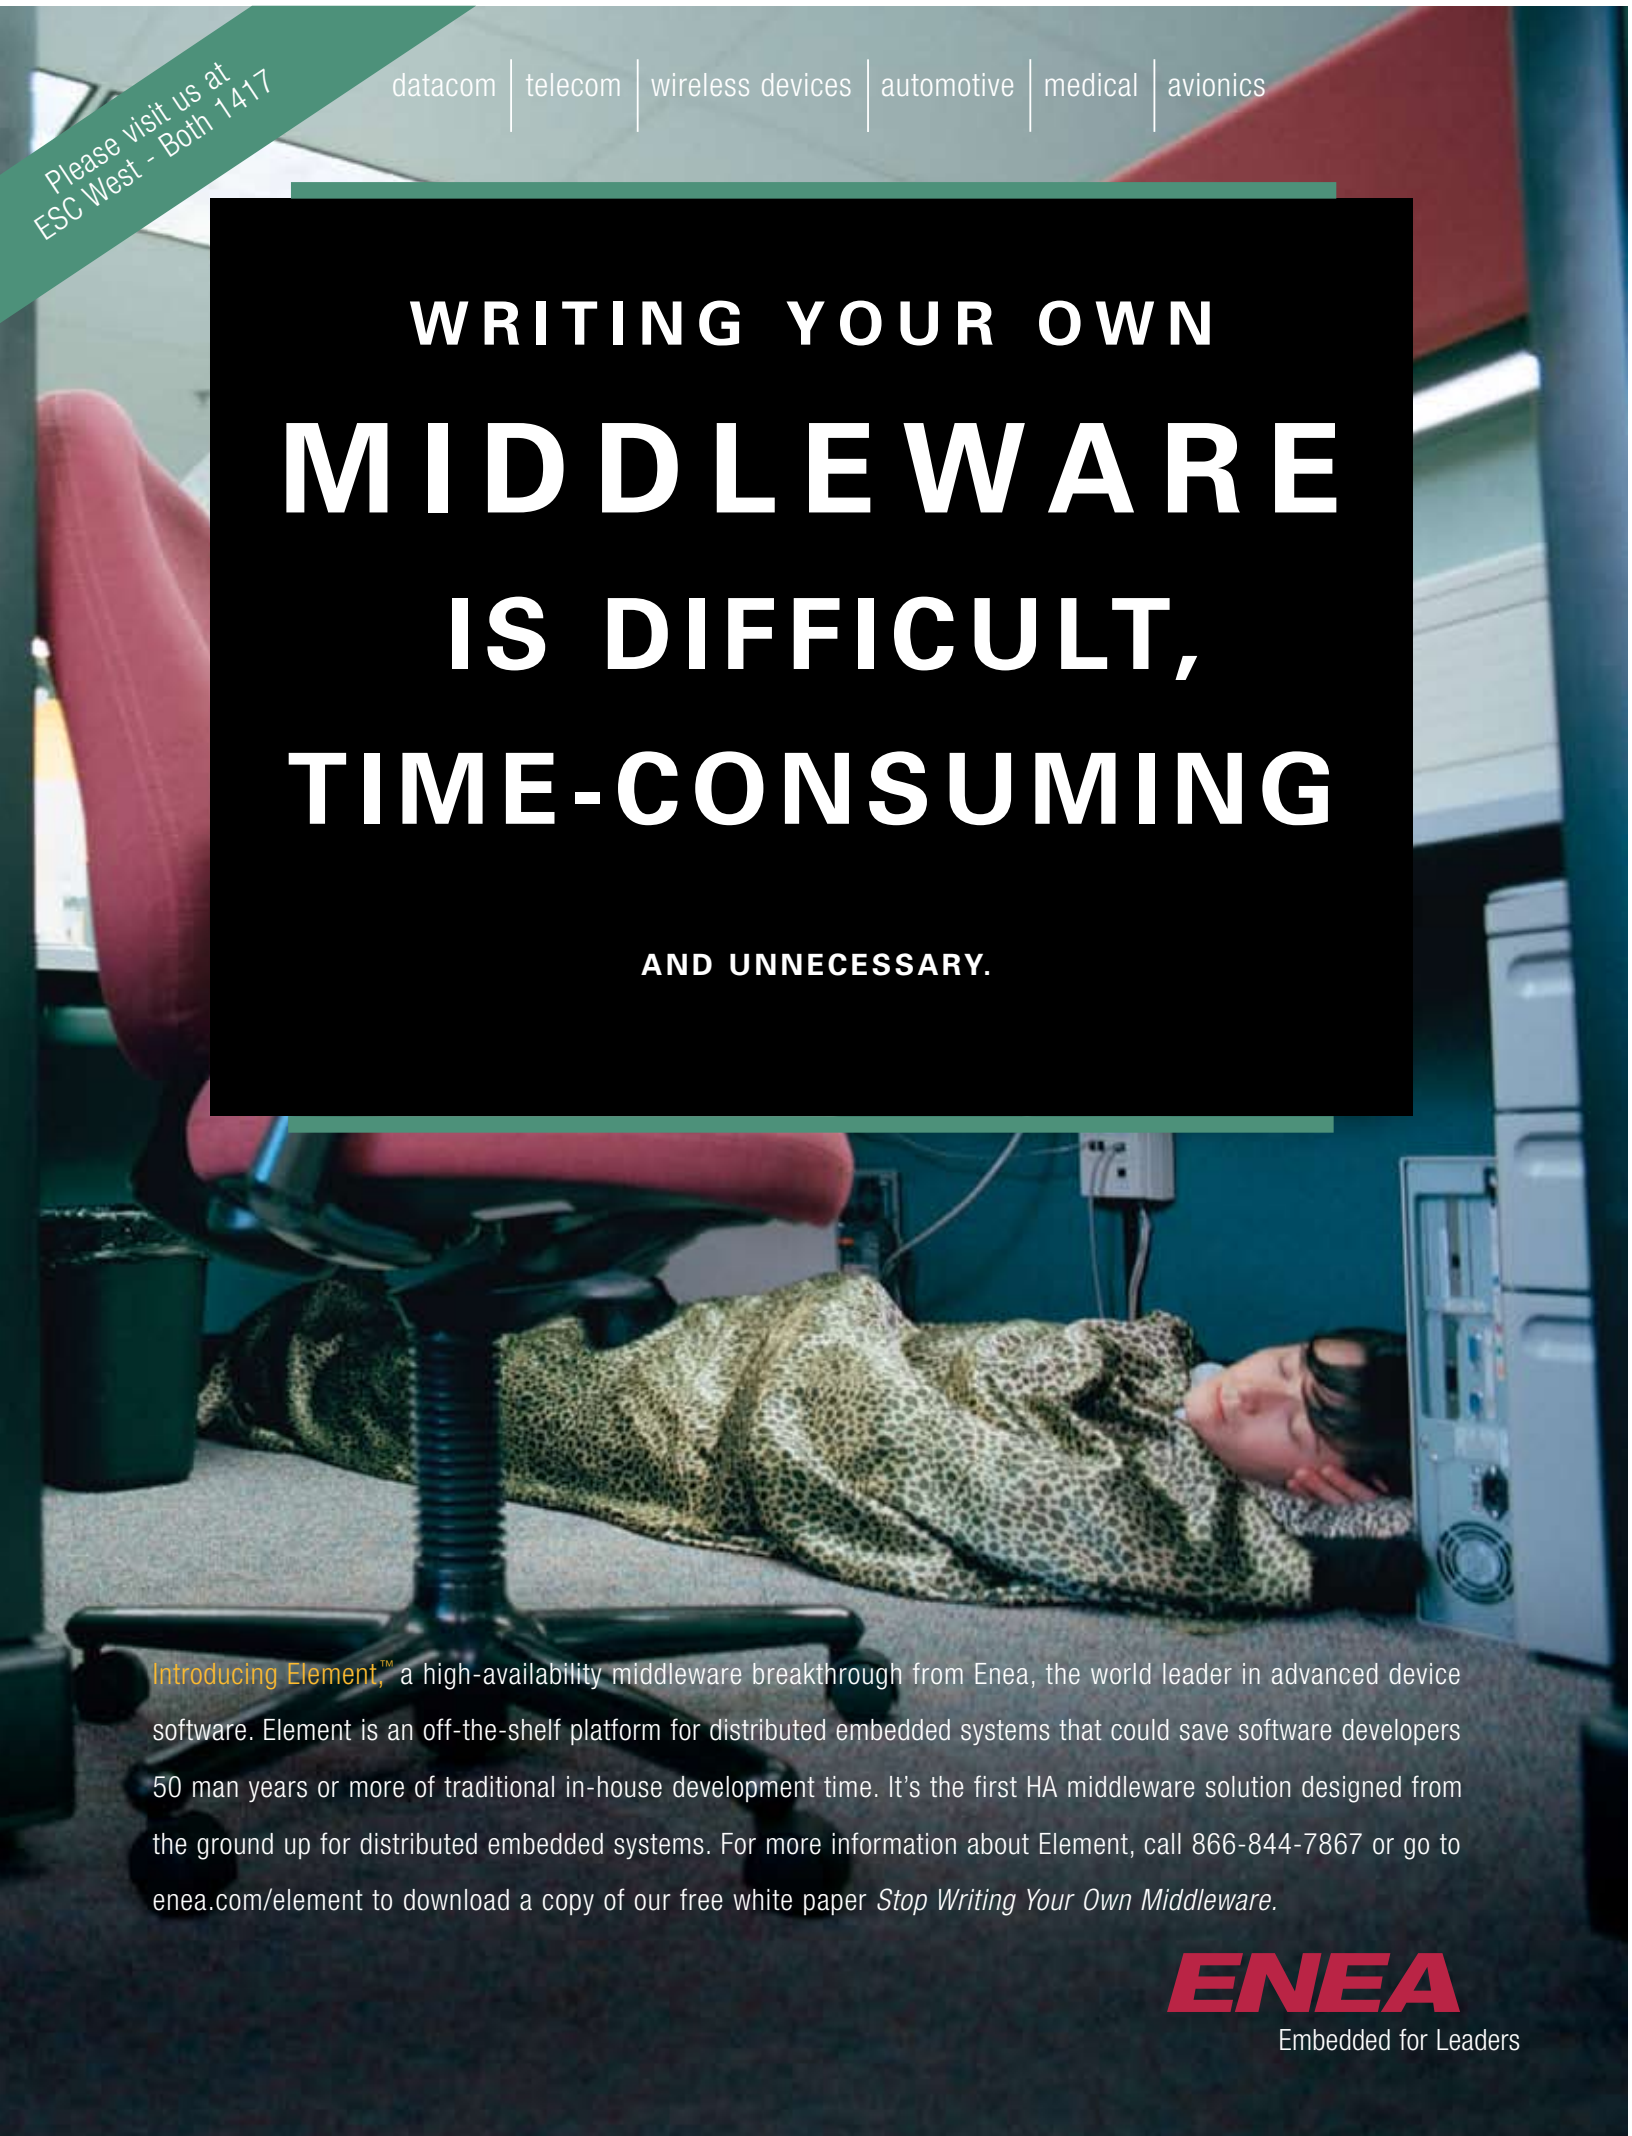
Please visit us at
ESC West - Both 1417
datacom | telecom | wireless devices | automotive | medical | avionics
WRITING YOUR OWN
MIDDLEWARE
IS DIFFICULT,
TIME-CONSUMING
AND UNNECESSARY.
Introducing Element,™ a high-availability middleware breakthrough from Enea, the world leader in advanced device software. Element is an off-the-shelf platform for distributed embedded systems that could save software developers 50 man years or more of traditional in-house development time. It's the first HA middleware solution designed from the ground up for distributed embedded systems. For more information about Element, call 866-844-7867 or go to enea.com/element to download a copy of our free white paper *Stop Writing Your Own Middleware.*
ENEA
Embedded for Leaders

EtherCAT: A Real-Time Industrial Fieldbus on Ethernet

EtherCAT fieldbus systems are based on technology and techniques designed to overcome the performance and determinism limitations associated with other attempts to utilize Ethernet packets and hardware to support an industrial fieldbus.

by Paul Fischer
TenAsys

A number of Ethernet fieldbus solutions use fieldbus masters to send Ethernet packets to individual fieldbus slaves, or broadcast to a group of fieldbus slaves. This usually takes the form of a fieldbus packet encapsulated inside an IP packet such as UDP. The fieldbus master must then wait for the slaves to process the received data and respond to the fieldbus master via a response packet.

In an EtherCAT fieldbus system, Ethernet packets are delayed only a few nanoseconds by each slave in the chain because the slave reads and writes to the encapsulated datagram each time it is received. Ethernet packets are "processed on-the-fly." Each slave extracts data bits inside the Ethernet packet addressed specifically to it, while simultaneously forwarding the packet to the next device in the chain. Similarly, data to be returned to the EtherCAT master is inserted on-the-fly while the packet is processed by the slave, as the Ethernet frame passes through the device.

EtherCAT fieldbus masters are built using commercially available standard network interface cards (NICs). EtherCAT fieldbus slaves utilize a device called a Fieldbus Memory Management Unit (FMMU) that can read, write and forward EtherCAT Ethernet packets on-the-fly. It is possible to build high-performance, deterministic, fieldbus systems that are flexible, expandable and cost-effective using standard Ethernet cables and connectors.

Ethernet Header			ECAT	EtherCAT Datagram			Enet
DA	SA	Type	Frame HDR	EtherCAT HDR	Data	WKC	FCS
(6)	(6)	(2)	(2)	(10)	32...1486	(2)	(4)

Figure 1 A single EtherCAT data object delivered inside a single Ethernet packet.

An EtherCAT fieldbus avoids the determinism problems associated with packet collision on standard Ethernet networks (like most TCP/IP networks in operation today) by isolating the fieldbus network segments from conventional network segments. Ethernet packets are placed onto the EtherCAT network by an EtherCAT master where they are processed on-the-fly by EtherCAT slaves. This process on-the-fly technique accommodates bandwidth utilization approaching 97% utilization, or the equivalent of 97 Mbits/s using 100 Mbit/s Ethernet hardware.

Standard Ethernet

A key goal of the EtherCAT fieldbus is to use standard Ethernet packets so that one can utilize standard, low-cost, 100 Mbit/s Ethernet connectors, cables, switches and silicon (for 100Base-TX and 100Base-FX media). Packing useful bits and bytes into each Ethernet packet for on-the-fly read and write operations by multiple EtherCAT slaves is the key to maximizing usage of standard Ethernet packets.

By definition, Ethernet packets designed for 100 Mbit/s media have a minimum size of 64 bytes and a maximum size of 1518 bytes. After subtracting the standard Ethernet packet overhead of 18 bytes (Destination Address, Source Address, Frame Type and Frame Checksum), a minimum size Ethernet packet can deliver a 46 byte payload, whereas a maximum size packet can deliver 1500 bytes of data. An EtherCAT datagram, the

Get Connected
with companies mentioned in this article.
www.rtcmagazine.com/getconnected

EtherCAT data object stored inside an Ethernet packet, has an overhead of two bytes per Ethernet packet plus 12 bytes per EtherCAT datagram. This leaves room for 32 EtherCAT data bytes in a minimum size Ethernet packet to 1486 EtherCAT data bytes in a maximum size Ethernet packet (Figure 1).

An Ethernet packet that contains an EtherCAT datagram can be easily identified as an EtherCAT packet because it contains an EtherType number in the Type field equal to 0x88A4 (hex 88A4), a number that has been issued by the IEEE Registration Authority for use with EtherCAT-type Ethernet packets. It also ensures that standard Ethernet hardware can transmit and receive EtherCAT network packets.

A single Ethernet frame can contain more than one EtherCAT datagram, which might be appropriate for complex EtherCAT fieldbus systems. In that case, the Ethernet packet takes on the format in Figure 2.

If an EtherCAT packet needs to be routed over a standard TCP/IP network, it can be transported inside a UDP packet, where it is marked for easy identification using UDP port number 34980 (0x88A4), a number assigned by the Internet Assigned Numbers Authority (IANA) for transport of EtherCAT packets over TCP/IP networks. EtherCAT packets embedded inside a UDP packet have a format similar to those shown in Figures 1 and 2.

Ethernet Header			ECAT	EtherCAT Datagram			EtherCAT Datagram				Enet
DA	SA	Type	Frame HDR	EtherCAT HDR	Data	WKC	EtherCAT HDR	Data	WKC		FCS
(6)	(6)	(2)	(2)	(10)	32...1442	(2)	(10)	32...1442	(2)		(4)

Figure 2 Multiple EtherCAT data objects delivered inside a single Ethernet packet.

Near Maximum Performance

By efficiently packing EtherCAT data into each Ethernet packet it is possible for an EtherCAT network to utilize nearly 100% of the available bandwidth of 100 Mbit/s. Achieving high bandwidth utilization, however, also requires that one step back from the standard collision detection (CSMA/CD) mechanism that is common to most Ethernet networks.

Ethernet packet collisions are avoided by controlling access to the EtherCAT fieldbus, or by segmenting the network into EtherCAT network segments. A typical EtherCAT master transmits packets onto an EtherCAT segment on a regular periodic schedule. Packets are forwarded by the EtherCAT slaves on a segment, which process them on-the-fly for both read and write operations. Packets are not stored and forwarded by EtherCAT slaves; rather, they are inspected and retransmitted on-the-fly as they pass through each EtherCAT slave. Standard Ethernet cabling is wired for full-duplex operation over the segment so that when a packet reaches the last slave in a segment it can be returned to the EtherCAT master, via the receive lines of the Ethernet cable, again passing through each EtherCAT slave in the network segment in reverse order.

The effective utilization of an Ethernet packet varies between 50% and 98%, depending on the size of the EtherCAT datagram, for those scenarios where one cycle of fieldbus data is transmitted inside a single EtherCAT datagram within a single Ethernet packet. As mentioned previously, a maximum size Ethernet packet, containing one EtherCAT datagram can hold 1486 data bytes, or 11,888 data bits.

Because no collision detection is required on an EtherCAT segment, we can send out packets at the maximum rate the cable will transmit. Adding 20 bytes of "time" for the inter-frame and preamble overhead required of standard Ethernet hardware, the theoretical percent utilization of the Ethernet cable, or bandwidth utilization, ranges from 38% for a minimum size packet up to 97% for a maximum size packet.

In other words, a theoretical throughput of 38 Mbits/s to 97 Mbits/s is possible

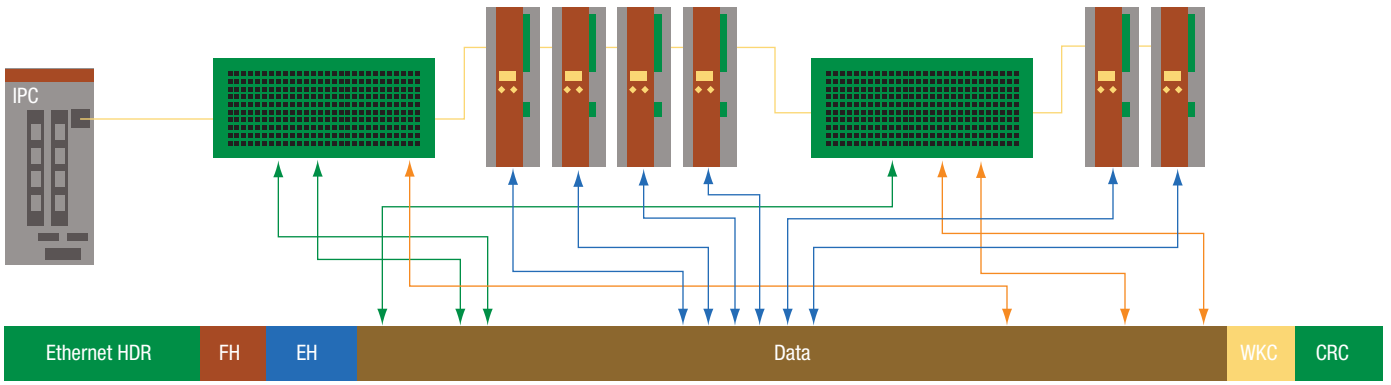

Figure 3 Logical addressing for data inside a single EtherCAT datagram.

Embedded Box PC
powerful, fanless affordable

MICROSPACE® PC Value

MPCV855
_ Intel® Processor 600 / 800MHz
_ 256 MB RAM, 40 GB HDD
_ 4x USB, 2xCOM
_ 2x LAN (1 GB/100 MB)
_ VGA and LVDS
_ expandable 2x PC/104-*Plus*, 1x MiniPCI

MPCV855I with expansion
_ LPT1, COM3/4 as RS232/422/485

MPCV855M with expansion
_ DVI-D and Sound

Further information:
www.digitallogic.com/mpcv_en

Embedded Systems Conference Silicon Valley
Booth 1538

DIGITAL-LOGIC offers a wide range of embedded computer systems for stationary and mobile applications

Advanced Digital Logic Inc.
4411 Morena Blvd. Suite 101
San Diego, CA 92117-4328
Phone +1 858 490 0590
Fax +1 858 490 0599
sales@adlogic-pc104.com
www.adlogic-pc104.com

DIGITAL-LOGIC
smart embedded computers

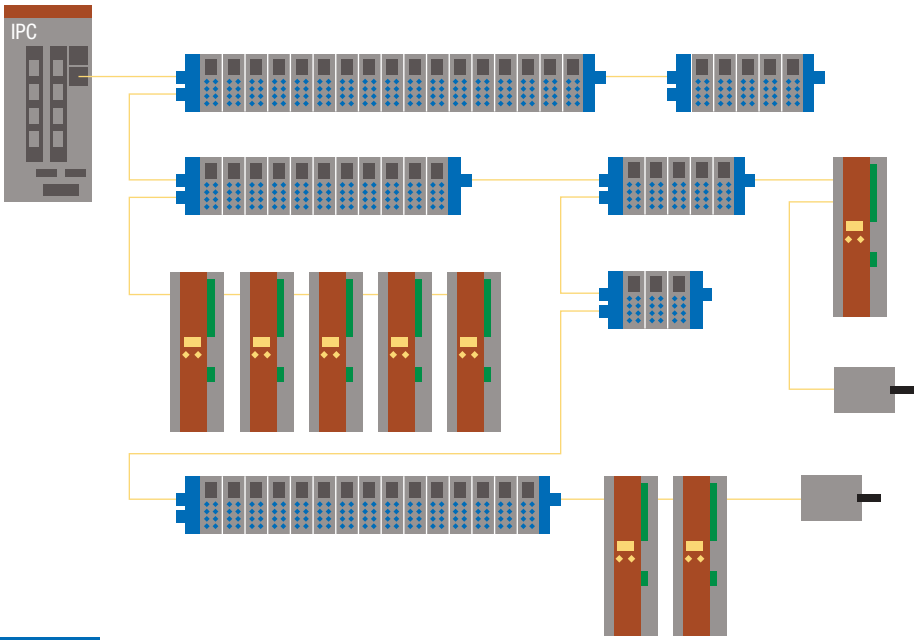

Figure 4 EtherCAT tree topology.

for an EtherCAT fieldbus segment. This assumes that all the data required of an EtherCAT segment fits within one EtherCAT datagram, which fits inside one Ethernet packet, and packets are placed on the cable at maximum throughput.

Efficiency in Addressing

EtherCAT datagrams can contain any mixture of reads and writes. Just as a single packet can address multiple devices, it is also possible to perform reads and writes within a single packet.

All the slave devices on an EtherCAT segment share a single MAC address. In effect, an EtherCAT segment can be thought of as one large Ethernet device. Each Ethernet packet is shared by all the EtherCAT slaves on a network segment. Ethernet frames are transmitted through the slaves on the segment, in sequence, from the first slave device (the owner of the MAC address of the segment, or the "Segment Address Slave") to the last.

Logical addressing, built into the EtherCAT protocol, is used to efficiently pack the data delivered by an Ethernet frame. Each slave addressed by an EtherCAT datagram is, in essence, reading from or writing to a shared memory segment stored on the EtherCAT master. That memory segment is being "shared" by the master with the slaves thru EtherCAT datagrams. Slaves read from the datagrams and write to the datagrams on-the-fly, as they pass through (Figure 3).

Each slave is assigned space in the master's shared memory block. This assignment is done at the bit level. A slave is assigned the following:

- A bit-oriented, starting logical address to locate its data in an EtherCAT datagram
- A bit-oriented physical address within the slave device that corresponds to the logical address assignment
- A bit length consistent with the physical address type (from 1 to 64)
- A direction that specifies if this is a read or write location (input or output)

More than one of these logical-to-physical assignments may exist for each slave. For example, a simple digital I/O device, with one input bit and one output bit would have two such assignments. Each assignment would be of length 1. One logical address would be assigned as an input bit and the other would be assigned as an output bit. More complex devices would require more bits, depending on their physical function and internal register requirements.

Determinism at the Master

EtherCAT devices are capable of processing Ethernet packets at a very high speed with loop rates ranging between 10 and 90 kHz. Obviously, then, the actual loop rate achieved will depend heavily on the ability of the EtherCAT master to

transmit and receive data over the Ethernet cable, while simultaneously processing that data in its internal process image (the shared memory segment that contains a reflection of the fieldbus device registers).

Virtually all modern NICs include a DMA engine, meaning that the process of transferring data from the process image to the fieldbus interface card is eliminated, requiring no CPU overhead. Instead, the NIC's DMA engine can transfer data directly from the process image to Ethernet packets on the cable. Conventional fieldbus interfaces require that the control computer CPU copy data from the process image to the interface card's internal RAM.

Maintaining synchronism between the process image and the EtherCAT packets then becomes important. Determinism at the EtherCAT master is a desirable trait for achieving stable loop rates. Because EtherCAT masters can utilize standard NIC hardware, common industrial PCs or specialized embedded PCs can be used without concern for the availability of specialized fieldbus adapter hardware. Real-time operating systems, Linux, and even Windows can be used as the operating system to drive the EtherCAT master. In the case of a general-purpose OS, such as Windows, it is useful to host the EtherCAT master on a real-time platform that works in conjunction with Windows.

Not Just About Performance

EtherCAT was not just designed to be the highest performing fieldbus that uses standard Ethernet packets, it also includes features that are an absolute necessity for real factory-floor applications. Features like: distributed clocks for synchronization, flexible cabling topologies, network redundancy, support for safety protocols and built-in diagnostics.

A clock jitter of significantly less then 1 microsecond is achievable within an EtherCAT segment. Because each EtherCAT segment is organized as a ring, every device in a segment can determine its "time position" within the ring by measuring the time between transmission and return of an EtherCAT packet. This propagation delay data can then be used to adjust the distributed clocks of each slave device so that they will run synchronously. EtherCAT segments support a variety of topologies, such as line, tree and star.

While a star technology is most commonly applied to switched office networks, it is not likely to be used by most EtherCAT networks. The line and tree topologies are more conducive to fieldbus applications because they require fewer connections and utilize a much simpler and more flexible cabling schema. Additionally, switches and hubs are not required by EtherCAT segments configured as lines or trees (Figure 4).

Optional cable redundancy enables devices to be changed (hot plugged) without having to shut down the network. Redundancy is very inexpensive: requiring only an additional Ethernet interface in the EtherCAT master and a return cable to turn a line topology into a ring topology. Switchover takes one cycle. EtherCAT also supports redundant masters with hot standby. An EtherCAT slave immediately returns a frame if an interruption occurs, so that failure of a device does not lead to the complete network being shut down.

EtherCAT networks are designed to simultaneously support the transport of safety and control communication over the same network. The EtherCAT safety protocol is certified to IEC 61508 and meets the requirements of Safety Integrated Level (SIL) 4. EtherCAT safety data can be routed without a need for safety routers or gateways.

Bit faults during the transfer of packets are reliably detected by evaluating the Ethernet packet's CRC checksum. Automatic monitoring of error counters by slave devices allows for precisely locating defective network sections. Intermittent errors, such as those caused by EMI, bad connectors or damaged cables can also be detected and located. ◢

TenAsys
Beaverton, OR.
(503) 748-4720.
[www.tenasys.com].

EtherCAT Technology Group
[www.ethercat.org].

Analogue & Micro Ltd

PowerPC Embedded Modules and Development Boards

Ready out of the box!

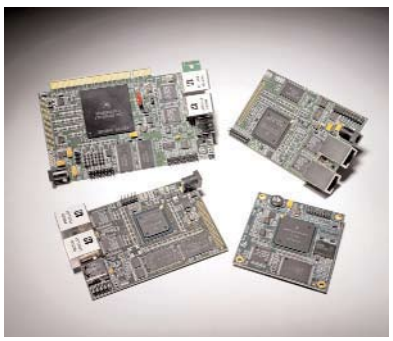

PowerQuicc™ embedded modules & development platforms available with Linux and eCos operating systems. With the ported Linux and eCos operating systems, programming examples, supplied cables and power supply, the products are *ready to go out of the box!*

Modules and boards are available for the MPC850, MPC852T, MPC866T, MPC870, MPC875, MPC880, MPC885, MPC8250VR, MPC8270VR, MPC8280, MPC8247, MPC8248, MPC8271, MPC8272, MPC8343, MPC8347, MPC8349, MPC8541, MPC8555, & MPC5200.

USI Ultimate Solutions, Inc.

10 Clever Lane
Tewksbury, MA 01876 USA
Toll Free: 866.455.3383
Fax: 978.926.3091
Email: info@ultsol.com
Web: www.ultsol.com

INtime® keeps Windows® on time.

Now the world's most desirable human interface can have time-critical responsiveness. Windows XP and our INtime RTOS run on the same platform, so no application is too fast for Windows. You'll cut through timing jams and keep critical events under control. INtime's robust stability and time-tested reliability handle hard real-time applications easily. The beauty of systems using INtime 3.0 is keeping all the advantages of a Windows environment without adding hardware cost or complexity.

Speed your design cycle with power and grace. Using Microsoft® Visual Studio® tools to develop and debug both Windows and INtime applications is easy and saves time. You'll pass the competition with real authority.

It pays to be prompt. Learn how the union of Windows and INtime really works, and how 25 years of practical innovation profits you. Call toll-free (877) 277-9189 or visit www.tenasys.com/intime

TenAsys®

INtime 25 years

Intel® Communications Alliance Affiliate Member

Microsoft® Windows® Embedded Partner 2004 Partner of the Year Independent Software Vendor

Microsoft® Windows® Embedded Partner 2005 Partner of the Year Independent Software Vendor

Optimized for Visual Studio

Copyright © 2006 TenAsys Corporation. All rights reserved. TENASYS, INTIME, and IRMX are registered trademarks of TenAsys Corporation. Other trademarks and brand names are the property of their respective owners.

DDS Over Ethernet Simplifies Industrial Automation

Industrial automation has used a client-server structure that is proving to be a barrier to maintenance, configurability and real-time performance. DDS can eliminate the barriers of a client-server architecture, while utilizing Ethernet connections as the backbone for many industrial systems.

by Bob Kindel
Real-Time Innovations

Industrial automation (IA) typically has many elements gathering information and feeding that data to other elements that make decisions and control processes. The traditional approach to IA uses a client-server architecture. All the elements in the system connect to a central information server that collects data from the information-gathering elements, makes decisions and issues commands to the process control elements. Variations on the architecture allow the process control elements to make decisions on their own, leaving the central server to act as a clearinghouse for collecting and distributing data. Mixed systems, with some control elements autonomous and others slaved to the server, are also common.

There are several problems with these approaches. One is that they have a single point of failure: the server. If it goes down the entire system collapses. Building failover redundancy is both difficult and costly in a client-server architecture. The central server also serves as a data bottleneck in the system because it must handle each piece of data twice: once to receive and once to send.

Another problem with the client-server architecture is that it becomes increasingly difficult to modify, upgrade and maintain as the system becomes more complex. In the least flexible configurations, the signal paths between elements are hardwired to fixed portals. Such systems can be brought to a halt by simply plugging Sensor A into the port designated for Sensor B. More flexible systems use a network connection, such as Ethernet, to link elements. While networks make the physical connections more forgiving, the logical connections between elements can still be 'hard-wired." Adding a new

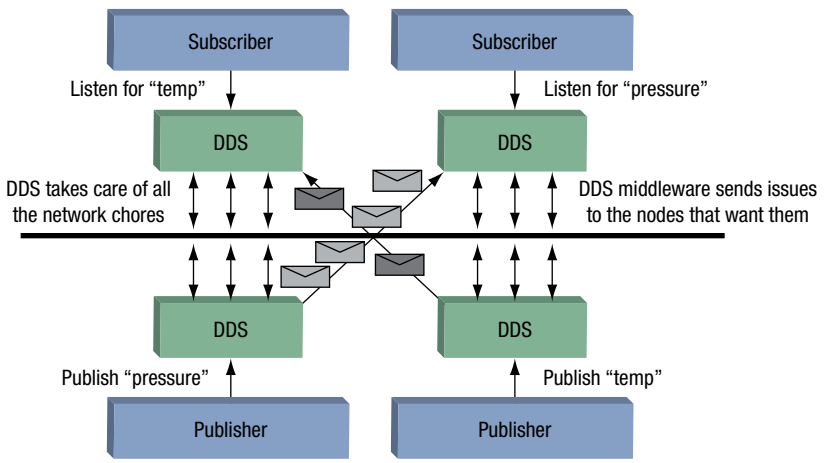

Figure 1 Industrial automation applications can result in complex arrangements of sensors and control points that become increasingly difficult to manage as they expand if they use the traditional client-server architecture.

Get Connected
with companies mentioned in this article.
www.rtcmagazine.com/getconnected

element to such a system requires reprogramming the server so that it will accept the new data and send it to the appropriate destinations.

The inflexibility of client-server architectures can cause many problems for industrial automation users. As systems become more complex, the effort to program them becomes increasingly difficult. Similarly, because units must be "wired" correctly to avoid system crashes, system maintenance becomes more difficult as well as riskier. Mistakes become harder to find and the high cost of system downtime makes mistakes more costly.

Client-server architectures also have the drawback that they lack determinism and can have a high latency. When a sensor element has data ready to send, it must wait for the server to stop being busy before the data moves into the system. Devices waiting for that data then face an additional delay waiting for the server to become available to pass the data along. Whether using a round-robin approach or being interrupt-driven, the server cannot guarantee fixed timing in the delivery of data and there is no way of synchronizing the availability of data with the needs of the recipients.

Figure 2 The ALTSOM Schilling Quest is an example of an autonomous underwater vehicle that can be configured for different missions. That reconfiguration is aided by a hierarchical software architecture where components interface to a single data communications scheme.

Alternatives to Client-Server

The industry has developed alternative approaches to overcome the limitations of client-server architectures in industrial automation. However, many of these alternatives have used proprietary architectures and interfaces, often leaving system developers at the mercy of a single supplier in order to obtain the sensors and control elements they need. This can become a crippling problem if the system needs an element that is not commercially available yet. Designing a custom device based on an outside party's proprietary hardware and software interface can be both difficult and expensive.

There are approaches that use industry-standard interfaces, however. One approach, developed by the Object Management Group (OMG) is the data distribution service (DDS). This standard specifies a middleware that sits between the application layer and the transport layer of a network's OSI stack. DDS can operate on top of both IP and non-IP transport mechanisms, including Ethernet, shared memory, InfiniBand and StarFabric. DDS follows a publish-subscribe model that eliminates many of the client-server architecture's drawbacks.

In the publish-subscribe model there is no central data server, hence no single point of failure. Instead, data flows directly from source to destination. Data sources put data onto the network (publish) as the data becomes available, tagging that data for receipt by all registered subscribers. Control elements that need data alert the data sources to their needs by registering with the data source as a subscriber. Once all the subscriptions are registered with the data sources, data flows automatically (Figure 1).

Adding a new element to a publish-subscribe network becomes trivial. A new data element announces itself to the network when it attaches, optionally describing the types and formats of the data it provides. Subscribers then identify themselves to the data source to complete the system reconfiguration. A similar process allows a new control element to connect with data sources when attached. Network configuration thus virtually reduces to a "plug-and-play" exercise.

DDS can be made to be a perfect match to the needs of industrial automation by providing a variety of data services to support a variety of needs, from reliable command and control to real-time data flow. For example, a STOP command to equipment can use a data service type that guarantees error-free delivery. On the other hand, time-sensitive data can use a service type that minimizes latency. A temperature sensor, for instance, could send a stream of data under DDS using a data service type that provides only a "best effort" rather than guaranteed delivery. This data service type supports recipients that only need a current reading, not a full history, but need that data as soon as it becomes available. In such cases, the overhead of retransmitting corrupted data to guarantee delivery would interfere with timely delivery of current data.

DDS Benefits

The DDS standard provides IA developers with numerous benefits. To begin with, the standard defines an applications programming interface that standardizes the way units request and present information. This allows the software of individual units to be constructed without requiring knowledge of other units in the system. The application simply defines the type and characteristics of the data it will handle. The lower layers of DDS handle the details of establishing the data pathways based on the devices that are present on the network. This degree of isolation also means that units do not need a common software operating system; the DDS allows even a collection of heterogeneous units to exchange data.

Because the publish-subscribe architecture does not require a central data server (although the network may require a network manager), system reliability is

relatively easy to ensure. A critical data or control element may be constructed to provide a failover response utilizing a redundant element without affecting the data flow. The redundant data element simply steps in to replace the failed unit, publishing data to the same subscriber list. A standby control element can subscribe to the same data as its primary element so that the two remain synchronized in case of failure.

Many IA systems exist that demonstrate the benefits of the publish-subscribe architecture that DDS offers. Applications as diverse as a semiconductor manufacturing plant, an experimental fusion reactor, and an underwater remotely operated vehicle (ROV) have gained design flexibility and eased system development by using NDDS 3.x over Ethernet, a real-time publish-subscribe middleware product that came before the DDS standard and helped shape its definition. The diverse IA systems using this publish-subscribe middleware all share a common theme: they contain complex mixes of data and control elements that need configuration flexibility.

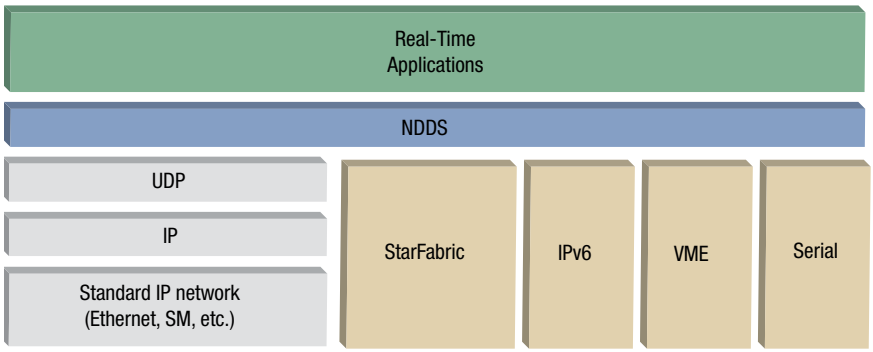

Figure 3 Publish-subscribe middleware sits between applications and the transport mechanisms in the OSI layer model.

Industry Examples

The Nikon Semiconductor wafer fabrication plant, for example, has an ongoing need for equipment upgrades while maintaining compatibility with the legacy elements. To address this need, Nikon developers used NDDS to send medium data-rate communications such as trajectories, commands, and state information between computers running different operating systems. The NDDS also served as a bridge between new and legacy equipment designs, and allowed equipment to maintain compatibility from one generation to the next.

The Max Planck Institute for Plasma Physics (IPP) created a new control system for its Tokamak nuclear fusion device using NDDS. The control system generates data at a rate of several Mbytes/s during experiments, and passes the data to a bank of video consoles and archival storage as well as feeding it back to the controllers to modify operating parameters during the experiment. The configuration of feedback paths and data streaming to the video consoles can vary from experiment to experiment, so the system continually needs reconfiguration. Using NDDS simplifies the reconfiguration process, eliminating the need for network programming and allowing experimenters to simply subscribe to the data sources they want to feed back into the control structure and want to monitor.

ALSTOM Schilling's Quest ROV is used for servicing, exploration, salvage and recovery operations (Figure 2). The network serves to link the ROV to the ship for control and communications. By using NDDS, the development team was able to design the system as a hierarchy of interacting components and then test each component and all higher-level assemblies independently. This hierarchical structure also simplifies the inevitable customization each customer requires as well as fostering design reuse. New components insert easily into the publish-subscribe model and component modifications have only minor impact on other system components.

PEG® the leading GUI for Embedded Systems

Prototype Today!! FREE EVALUATION PACKAGE

- PEG+ - Full Featured Windowing in C++
- C/PEG - Smallest Footprint in ANSI C
- Royalty Free
- Fast execution speed
- Completely ROM-able
- Delivered with Full Source Code
- Development Tools including FontCapture, PEG WindowBuilder, and ImageConvert
- Complete set of screen drivers included
- Completely customizable
- Industry leading RTOS Support
- Supports all popular target processors, video controllers and I/O devices
- Multi-lingual support – 2-byte character sets & UNICODE string encoding
- Event-driven programming model
- Application Design Services
- Knowledgeable and timely support to users around the globe
- Now includes a fully licensed version of Paint Shop Pro 9

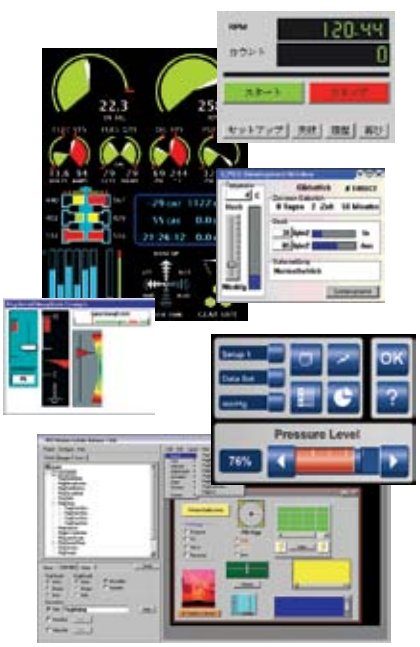

SWELL software

GRAPHICS SOFTWARE FOR EMBEDDED SYSTEMS

WWW.SWELLSOFTWARE.COM 810-982-5955

DDS Limitations

For all of its benefits, the DDS standard as defined by the Object Management Group does have some limitations. The standard focuses on the applications programming interface for DDS but does not specify the implementation. This means that developers of DDS middleware each create their own variations, making different versions somewhat incompatible. While they will all interoperate at the application layer, they may not be able to communicate on the same wire due to protocol and bit pattern differences. This incompatibility is just a temporary situation, however. Companies such as Real Time Innovations (RTI) are working with the Object Management Group to establish a standard protocol to allow interoperability. Adoption may be as little as a year away.

The DDS standard also leaves room for variations that can affect the performance and code size of a DDS implementation in a given application. For instance, the RTI implementation—NDDS—offers additional real-time features that are not covered by the standard. These features include the ability to use the user datagram protocol (UDP) instead of the transmission control protocol (TCP) to make data transfers more deterministic and to allow anonymous operation by system elements, making peer-to-peer transactions simpler and easier (Figure 3). A small-footprint version of NDDS that implements a subset of the quality of services options is also available, bringing the publish-subscribe model to even small control and data units.

Not all DDS implementations provide the low latency and the determinism needed for real-time control, or offer small-footprint versions. What they all offer, however, are compelling advantages for industrial automation. The isolation between the application layers of individual units that DDS creates simplifies system configuration and equipment upgrades while also reducing programming overhead. The publish-subscribe model eliminates the data server as a single point of failure and throughput bottleneck. Additionally, its standards base means that developers can focus more of their attention on their application software and less on trying to get a heterogeneous set of components to work together. ◢

Real-Time Innovations
Sunnyvale, CA.
(408) 200-4700.
[www.rti.com].

RICOH ELECTRONICS, INC.
PROVIDING MANUFACTURING & DESIGN SOLUTIONS FOR YOUR EMBEDDED SYSTEMS!

Ricoh Electronics, Inc. is a leader in manufacturing and developing high quality, long life embedded systems. We provide our customers with standard boxes, motherboards and chassis products for their embedded systems. These systems are manufactured in our vertically integrated manufacturing plant in Tustin, California. Our factories are ISO-9001 and ISO-14001 certified and RoHS compliant. We provide our customers with the following manufacturing services:

- ✔ Design for manufacturability
- ✔ Printed circuit board assembly
- ✔ Systems integrations and product assembly
- ✔ Metal and plastic parts design and fabrication
- ✔ Logistic support, including drop-shipping to your customers/end-users
- ✔ Supply chain management

Visit us at Embedded Systems Conference Show Booth #1941 for a FREE gift!

Ricoh Electronics, Inc. Embedded Products:

- ✔ Intel® embedded chipset architecture
- ✔ Customization of product and BIOS
- ✔ Qualified and tested board components

RICOH

RICOH ELECTRONICS INC.
One Ricoh Square
1100 Valencia Ave.
Tustin, CA 92780, USA
Phone: 714.566.6017
www.rei.ricoh.com

Headless, But Not Mindless.

Add Micro/sys brainpower to your embedded application.
There's a lot riding on that new OEM design: sophisticated application code, tricky I/O interfaces, and system integration issues. You don't need a headache with CPU design issues, too. You just want a reliable CPU that will plug into your application, interface with your system I/O, and run. Headless might be the right way to go, but the price has got to be right.

We have just the processing brainpower you need!
For example, the new Micro/sys SBC1586 is a headless PC/104 size Pentium® workhorse. It includes 10/100BASE-T Ethernet, USB, four COM ports and a removable CompactFlash® for data logging or OS. This embedded PC supports Linux, DOS, Windows® NT or CE, and VxWorks.® Starting under $500.

For a less scary ride on your development adventure, choose from a wide selection of X86 through Pentium CPUs, headless or fully configured with A/D, D/A, and other PC I/O. Call today and let the minds of our engineering and customer service teams help you rein-in those unbridled development costs by offering a compact and cost-effective solution.

MICRO/SYS

3730 Park Place, Montrose, CA 91020
Voice (818) 244-4600 Fax (818) 244-4246
info@embeddedsys.com

www.embeddedsys.com

For more fully-featured Micro/sys embedded PCs in various sizes, options and interfaces, take a look at the boards below and check our Web site for complete details:

SBC1586 (right)
Headless Low-Power Pentium PC/104 board with CompactFlash, 4 COM ports, USB and Ethernet. Starting under $500.

SBC1495 (middle)
PC/104 OEM 486/586 with VGA, COM, Ethernet, and CompactFlash with DOS, Windows, Linux OS.

SBC2590 (left)
EBX Pentium Controller with VGA, COM, Ethernet, A/D, D/A and your choice of Operating System.

©2004 Micro/sys. All brand and/or product names listed are registered trademarks of their respective owners.

PC/104 Modules: Stackable, Small, and Rugged

Begin your next design with WinSystems' PC/104 and PC/104-*Plus* modules. These compact, 90mm x 96mm (3.6" x 3.8") modules are easy-to-use and ideal for space-sensitive applications.

WinSystems' products include:

- ▶ Single Board Computers
  - Intel Pentium
  - AMD 520
- ▶ Input/Output (I/O)
  - Analog
  - Digital
  - Serial
  - GPS
  - USB 2.0
- ▶ Communications
  - Networking
  - GSM Cellular
  - 802.11 Wireless
  - 10/100/1000Mbps Ethernet
  - Modem
- ▶ Enclosures
- ▶ Long-life product availability
- ▶ -40°C to +85°C operation
- ▶ Quick Start SBC Developers kits for Windows® XP, CE, and Linux.

You'll find our PC/104 modules stack-up as the perfect fit for industrial, communication, transportation, medical, instrumentation, and MIL/COTS applications.

Call 817-274-7553 or
Visit www.winsystems.com

Ask about our 30-day product evaluation!

WinSystems®

715 Stadium Drive • Arlington, Texas 76011
Phone 817-274-7553 • FAX 817-548-1358
E-mail: info@winsystems.com

EPIC Express: A Bridge to the Future for High-Performance I/O

EPIC Express combines the performance advantages of PCI Express with the widely used PC/104 bus in a small, stackable format. The result is one of the broadest ranges of stackable, plug-on I/O available, as well as a clear migration path to future, high-performance embedded applications.

by Susan Wooley
Micro/sys

A technical evolution is propelling PCI Express to become the wave of the future for high-performance I/O in industrial/embedded applications. PCI Express provides a high-speed, high-performance, point-to-point link for interconnecting devices: data is transmitted on one set of signals and received on another set of signals. It therefore offers relief for data bottlenecks created by attempts to process an increasing amount of data in real time.

The new EPIC Express standard harnesses the performance advantage of PCI Express by defining a small, stackable format with high-performance connectors for the PCI Express interface. EPIC Express also incorporates the popular PC/104 I/O connectors. This gives designers access to one of the broadest ranges of "plug-on" or stackable industrial I/O available in the marketplace today.

Additionally, PCI Express bandwidth can be scaled upward to keep pace with the increasingly faster CPUs that will be developed in the next 10 years. The combination of these two buses on a single platform therefore makes the compact, rugged EPIC Express standard an ideal platform for today's and tomorrow's embedded systems. This is the kind of stable platform that is attractive to embedded system designers.

This is the exciting part of EPIC Express: it provides embedded system designers a clear roadmap for upgrading and migrating their systems toward the faster serial interfaces of the future, while allowing them to implement this migration at their own pace and in their own fashion. With PC/104 and PCI Express on one platform (Figure 1), system designers can map their applications' initial design and

EPIC Express Stacks
PCI Express and PC/104

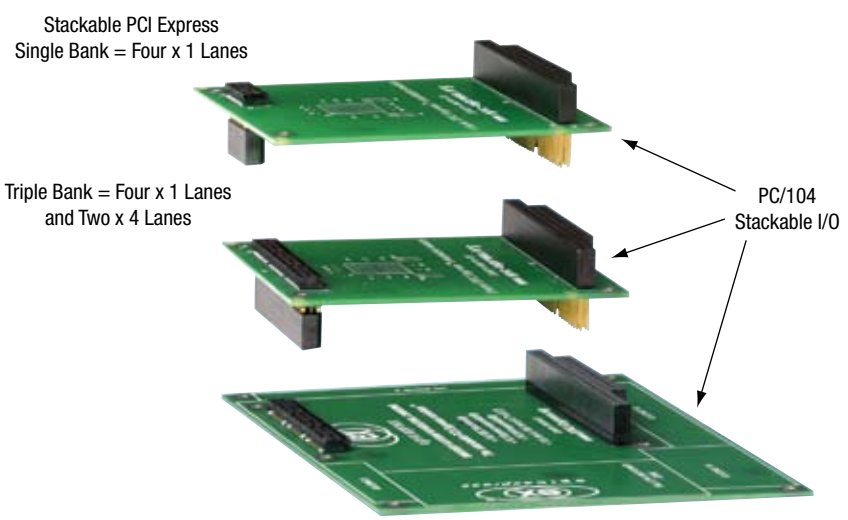

Figure 1 The new EPIC Express standard incorporates the performance advantage of PCI Express connectors as well as PC/104 I/O connectors in a small, stackable format. Including PC/104 and PCI Express on one platform gives a clean migration path for hardware and software.

Get Connected
with companies mentioned in this article.
www.rtcmagazine.com/getconnected

Figure 2 The AVATAR model remote control helicopter developed by USC's Autonomous Flying Vehicle Team utilizes a PC/104 stack for vision and another for flight control. It is capable of fully autonomous flight and includes flight-avoidance capabilities.

future upgrades by taking into account the future impact of emerging technologies. This makes a clean migration path from the standpoint of both hardware and software. It is also a key benefit of the EPIC Express specification: it provides a bridge to the future.

Adaptability, Scalability, Flexibility

EPIC Express places the two basic modes of communication—parallel and serial—onto one format. The PC/104 bus provides parallel communication as a simple I/O channel suited to industrial and control types of I/O, while PCI Express provides the high-bandwidth serial interface that moves data at Gbit/s rates suited to fast communication of large amounts of data. The combination of these two technologies on one format gives embedded system designers the ability to maximize their designs to fit specific size, power, environment or cost constraints.

Under the new EPIC Express standard, two configuration options are available for implementing PCI Express: Standard and Full. The Standard option supports four x1 lanes using a single 28-pin connector. The Full option supports four x1 lanes and two x4 lanes for future, higher bandwidth applications, as well as additional DC power.

With the Standard configuration, designers of applications that are cost-sensitive or have lower bandwidth requirements

Scratch
Standard Elma Type 32 Enclosure
Customized Elma Type 32 Enclosure
You don't need to start from scratch
Why start from square one when you can customize your enclosure from Elma's standard line of modular platforms? Use Elma's 40 years of expertise to build your exact configuration quicker, better, and more cost-effectively.
Elma has the largest inventory of standard parts and components in the U.S. And a knowledgeable technical support specialist will help you pick out exactly what you need.
So if you're looking for attractive quality enclosures but have a tight deadline, come to Elma. Our standard - or semi-custom - enclosures are right at your fingertips. No waiting! No tool-up! No kidding!
ELMA
Your Solution Partner
Phone: 510.656.0606 Fax: 510.656.8008 E-mail: sales@elma.com Web: www.elma.com
Handles
Industry's largest selection of handle styles
Classic, ergonomic, telco and long-telco styles
Strongest latching button
Patented alignment pin
Panels
Blue LEDs meet IEEE 1101.10 hot-swap specifications
Surface mount LEDs designed for IEEE front panels
Switches and test jacks designed for IEEE front panels

can implement PCI Express with a single-bank connector with four x1 links in a very small, cost-effective space. This is ideal for applications that need to take advantage of Gigabit Ethernet.

The Full configuration, with four x1 links and two x4 links, provides higher bandwidth, requires the full three-bank connector and has more power. This ensures that high-power, bandwidth-hungry devices, such as disk controllers and multi-gigabit communications devices, will perform at maximum throughput.

For now, EPIC Express does not implement x8 or x16 links. The exclusion of x8 reflects some current industry thinking that x8 PCI Express devices will be rare. However, the need is already apparent for a x16 definition to allow EPIC Express to take advantage of the full range of video cards in the future.

As a bridge standard, PC/104 is an extremely important part of the EPIC Express standard. It will probably be needed in most applications for some time to come. As the market matures, the convenience of a stable PCI Express standard will attract users of its own. How this will affect PC/104 on EPIC Express will depend on the success of the solo adoption of PCI Express, and on the market forces affecting production of the silicon required to implement PC/104.

Typical Embedded System I/O Checklist		EPIC Express Platform	
I/O Type	Selection Criteria	High Performance PCI Express	Standard PC/104
x86 processor	• Speed • Math coprocessor • Extended Temperature	✓	
Serial ports	• Type (16450 or 16550) • Levels (RS-232 or RS-422/ RS-485)		✓
Printer ports	• Will it be used for a printer? • Will it be used for digital I/O?		✓
Ethernet connection	• Speed: 10/100/1000	✓	✓
Video	• Display type – CRT, Flat Panel: LCD or TFT • Resolution	✓	✓
Serial Buses	• USB Type 1 or 2 • PCI Express • SATA	✓	
Operating Systems	• Windows Based • Linux • RTOS	✓	
Digital I/O	• Voltage level • Current source and sink capability • Initialization state		✓
Analog I/O	• Resolution • Range • Input type (single-ended or differential) • Conversion time • Method of starting conversions (hardware or software)		✓
Solid-state disks	• Type – Compact Flash, USB, DiskOnChip • Capacity • Physical size • Removability		✓

Table 1 EPIC Express's combination of PC/104 and PCI Express technology in the same platform helps designers maximize system performance based on the type of interfaces they need.

Support for Current PC/104 I/O and Expandable PCI Express Bandwidth

In the case of embedded systems that use the PC/104 bus for both an I/O control channel and a sophisticated user interface, continual improvements in video and communications can keep the system in an almost continuous state of re-design.

One example of this can be found in a project conducted by the University of Southern California's Autonomous Flying Vehicle Team. For several years, the students have been working on a model remote control helicopter that utilizes two PC/104 stacks: one for vision and the second for flight control. Named AVATAR (Figure 2), the helicopter is capable of fully autonomous flight and has recently been upgraded to include flight-avoidance capabilities.

AVATAR flies without human intervention and performs GPS waypoint navigation, autonomous vision-based landing and autonomous sensor-based takeoff. PC/104 technology has given the team a lightweight, low-power and high-performance system through two PC/104 stacks. One is a control stack that includes serial ports, timer/counters, sensor control and digital I/O. The other is a vision stack that captures images from three cameras.

As the USC Autonomous Flying Vehicle Team researches new capabilities—such as autonomous landing on a moving target, deployment on a moving target, stealthy target pursuit and vision-based obstacle avoidance in 3D—EPIC Express presents an exciting system platform for implementation of their goals. It can add high-bandwidth serial communication for camera systems and video requirements, which will allow AVATAR to perform even more complicated tasks.

A different type of application that can take advantage of EPIC Express is automated cleaning equipment for the metallurgy industry. The cleaning process is highly automated, but recently enacted industry quality standards now require documented microscopic snapshots verifying levels of cleanliness.

During the existing cleaning process, the cleaning solution is monitored for temperature, levels of contamination and time. These traditional electronic control processes are implemented with stacking PC/104 I/O boards. The sampling rates and time responses required for switching these controls are well suited to a parallel bus such as PC/104. However, upgrading the system with a video frame grabber to visually document the cleaning might be better implemented through a serial PCI Express expansion.

EPIC Express easily enables this mixing and matching of PC/104 and PCI Express

technology, which helps system designers maximize their system's performance based on the type of interfaces they need. Selection of the I/O platform is often determined by the general characteristics of the I/O type and its suitability for a particular embedded system. Selection criteria make some choices clear cut, while others are less obvious and require more system-level engineering (Table 1).

The Future of PC/104 and PCI Express

In the future, the data acquisition industry will be making advances in faster and higher-performance data gathering technology. As a stackable I/O channel, EPIC Express is an ideal platform for implementing this technology. In addition to its stackability, faster performance and compact size, EPIC Express offers system-level features made available by PCI Express. For industrial applications, these features include a software level to implement improved error detection, enhanced reliability, power management, hot-plug and free-flowing data. These higher-level system features will be required in the industrial controls of tomorrow to ensure the quality demanded in high-precision manufacturing environments.

EPIC Express brings greatly enhanced capabilities to system designers and is well matched to the demanding control requirements that emerging technologies will present during the next decade. Designers of control systems will seek to incorporate the technology being gained from military research and development into new tools for intelligence gathering equipment and weapon identification applications. Recent government funding for research and development in the emerging area of nanometer-scale technology will also result in the need for new manufacturing technologies and embedded control systems to implement emerging processes. The technologies being developed in these areas are finding their way into high-speed manufacturing lines, security applications, medical devices and automation processes.

In the pharmaceutical and chemical industries, designers are pursuing technologies to use in high-precision manufacturing processes for advanced drug delivery systems, including implantable devices, in the next few years. In medical applications, diagnostic tools such as cancer-tagging mechanisms and real-time diagnostics for physicians will need these enhanced control and monitoring capabilities.

In the field of wear-resistant coatings, the Department of Defense recently supported partnerships among the Navy, academia and industry to develop processes suitable for manufacturing more durable coatings for use in the marine environment. The result has been a nanoscale coating for use on air valves that will result in a $20 million reduction in maintenance costs over 10 years. The processes for manufacturing these coatings have required sophisticated embedded systems with control capabilities and user interfaces.

These advancements, as well as those not yet dreamed of, come with associated challenging system design issues and typical tradeoffs among heat, power, performance, size, design ruggedness and cost. This creates a future in high-performance embedded applications that is ripe for the EPIC Express format. ◢

Micro/sys
Montrose, CA.
(818) 244-4600.
[www.embeddedsys.com].

Fanless Super-computing!
An industry standard motherboard integrated with high performance Intel® Pentium® M technology...
ESC
Arcom Control Systems
APOLLO 1U ICE
APOLLO ICE 1U 19 inch rack mount industrial PC
• Pentium M CPU up to 2.1GHz
• Standard I/O connections from rear panel
• Front panel access to USB & IEEE1394 ports
• PCI riser card with 2 card expansion slots
• Front panel LCD status display and user LEDs
• ON/OFF/RESET switch with integral power LED
• Ideal for fanless, silent PC systems
• Just 14" (350mm) deep
• Hard disk activity indicator
• Tamper detect switch for enclosure lid
• Floppy disk, hard disk drive and CD/DVD drive
• AC ATX PSU: Auto-ranging 90 - 264V AC
• Support for Windows XP Embedded and Linux
APOLLO
• Intel® Pentium® M or Intel® Celeron® M processors
• 600MHz to 2.1GHz
• Up to 1Gbyte DDR memory
• Intel Extreme Graphics 2 dual video outputs (CRT & LVDS)
• Dual Ethernet - 10/100baseTx standard (option for 1000baseT)
• Six USB 2.0 channels
• Two IEEE1394a-2000 Firewire Ports (100/200/400Mbps)
• PCIbus and CompactFlash (CF+) expansion
• Industry standard EBX board format
APOLLO STANDARD
Arcom
www.arcom.com
888-941-2224
Think Embedded.
Think Arcom.

THINK INSIDE THE BOX.

EPIC ReadyBoard™ 800

PC/104 CoreModule™ 800

Mini-ITX MightyBoard™ 821

ReadySystem™ Turnkey Solutions

You know Ampro for building the finest embedded boards for transportation, military, industrial and medical electronics applications. But did you know we also build systems that complement many of the form factors we make? Our customers experience the fastest deployment of applications with Ampro's Industrial Computers.

READY-TO-GO OUT OF THE BOX...

because of the thinking that went into the box.

AMPRO COMPUTERS AND BOARDS

See us at ESC Booth 1610

Custom And Semi-Custom Systems Also Available.

AMPRO®

Call today (408) 360-0200 or visit: www.ampro.com

Ampro is a registered trademark of Ampro Computers, Inc. All other trademarks are the property of their respective owners. © 2006 Ampro Computers, Inc. All rights reserved.

SHARE YOUR VISION.

WE'LL SHARE OUR INNOVATION.

Share your vision with us. We'll customize our products to your requirements or partner with you to develop custom products all the way through deployment. Either way, you'll get leading-edge board level solutions for the most demanding embedded applications. Tell us what you need at **info@vadatech.com**.

- Comprehensive Hardware/Software solution for Intelligent Peripheral Management Interface (IPMI) version 2.0
- Symmetric Multi-Processing CPU for AMC/VME Modules
- A complete line of ATCA and AMC carriers
- High dynamic range A/D Converters

vadatech inc.
THE POWER OF VISION

www.vadatech.com 702.896.3337

PCI Express Stacks Up in High-End Embedded Applications

The new EPIC Express standard provides a unique module-stacking architecture for PCI Express that boosts the performance of PC/104-compatible applications.

by Colin McCracken
Ampro Computers

Years ago, the x86 personal computer got an I/O performance boost when the ISA bus width doubled from 8 bits to 16 bits. The first-generation bus architecture transitioned from CGA to VGA graphics and even 10 Base2 and 10 BaseT Ethernet. Bus bandwidth approached 16 Mbits/s.

In the embedded PC world, PC/104 then emerged as a stackable I/O standard, bringing this breakneck 16-bit ISA bus speed to small form-factor embedded applications. The simplicity of the interface attracted hundreds of new I/O modules, from A/D to networking to flat panel LCD to PCMCIA, facilitating the widespread adoption of PC/104 products in medical, transportation and industrial systems.

Back in the desktop world, bus bandwidth eventually became saturated with traffic, and cries for help were answered by the PCI bus. A wider and faster bus than ISA, its 32 bits per cycle time and 33 MHz meant that a gigabit of data could be transferred per second during bursts.

Second Generation I/O: PC/104-Plus

The embedded PC market faced the challenge of tracking the desktop world's bus migration, caused by an insatiable appetite for performance in the computer market. The PC/104-Plus standard answered the call by bringing second-generation bus architecture into the stackable card format. PCI bus expansion was complementary to the PC/104 (ISA) bus, and all of this bandwidth was still contained within the tiny 3.6-in. x 3.8-in. form-factor without sacrificing legacy compatibility.

Solid-state storage, whose reliability is required in many applications, got a performance boost in the form of 32-bit PC Cards, while 10/100 Ethernet, PCI graphics and 802.11 wireless technology began to take hold in the PC/104 market.

PC/104 Encounters a Paradigm Shift

PCI Express, formerly known as third-generation I/O (3GIO), is the rightful heir to the desktop motherboard backplane. A quick visit to the local computer and electronics store confirms this. The AGP graphics card slot has given way to a x16 PCI Express graphics slot. Several x1 slots occupy the next motherboard slot positions.

Each x1 lane brings a blistering 2.5 Gbits/s in transmit and receive directions simultaneously, full-duplex style. Even if data is moving in only one direction, a speed two-and-a-half times that of the parallel PCI bus renders PCI nearly obsolete from a bandwidth-per-pin perspective.

This bandwidth boost is accomplished with only five signals: Tx+ and Tx- for the transmit differential pair, Rx+ and Rx- for the receive pair and a clock line that generally is not used but is provided by the PCI Express specification in case devices wish to use it to generate their timing.

One or more token PCI slots round out the bus backplane. ISA is gone, replaced by a low pin count (LPC) bus primarily for the BIOS flash. PCI may have only a few chipset generations left before it disappears. PCI Express is the future.

PCI Express is the latest example of using high-speed serial technology to reduce size and cost. Like Ethernet and USB, PCI Express uses differential pair signals—such as Tx+ and Tx- for the transmit pair—to allow faster edge rates and data rates, so much so that the legacy parallel buses burdened with clock-to-signal skew cannot keep pace.

However, these Gigahertz-range signals require very specific trace widths, spacing and Z0 transmission line imped-

Get Connected
with companies mentioned in this article.
www.rtcmagazine.com/getconnected

Microsoft

© 2006 Microsoft Corporation. All rights reserved.

Microsoft, Windows, and the Windows logo are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. The names of actual companies and products mentioned herein may be the trademarks of their respective owners.

Apples and Oranges Working Together... That's the Power of Windows Embedded

The developers at Fujitsu know that the hundreds of peripheral drivers included in the Windows® CE operating system provide the solutions they need to easily integrate multiple device functions.

Needing to incorporate drivers for an infrared receiver, 802.11g WiFi card, and touch screen, the Fujitsu U-Scan Shopper development team chose Windows CE. Because Windows CE offers familiar, yet easily customized components and tools, they were able to develop the device in only five months. With the power of Windows CE, the Fujitsu U-Scan Shopper offers users coupons on demand, infrared scanners for product promotions, and in-store order placement, all from the convenience of the produce section, or wherever your shopping cart takes you.

"We considered Linux, but couldn't have achieved the same results, so we chose the Windows CE operating system."

— VERNON SLACK / **Store Solutions** / Fujitsu Transaction Solutions

The Power to Build Great Devices—get it with **Windows CE, Windows XP Embedded,** or **Windows Embedded for Point of Service.**

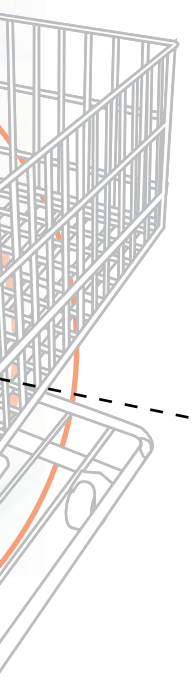

www.learnaboutembedded.com/shopping2

Windows Embedded

ances. Such requirements do not lend themselves well to pin-in-socket technology stacking connectors.

PC/104 has finally met its match. Or has it?

Stackable I/O Gets a Roadmap

Just when it appeared that thousands of system integrators and manufacturers using PC/104 off-the-shelf technology would face an uncertain future, including complete board-level redesigns, a group of five SBC technology leaders created a bridge to the future.

The EPIC Express specification defines PCI Express bus expansion for small form-factor SBCs using a unique stacking scheme. The specification even preserves legacy ISA (PC/104) bus support, due to its widespread use in embedded applications.

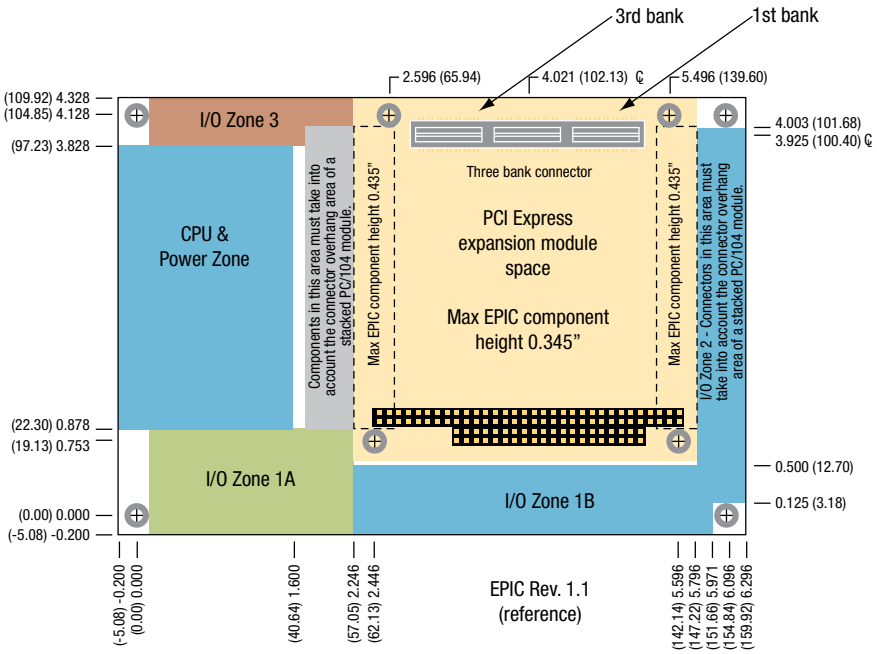

Figure 1 The EPIC Express form-factor defines a 1-bank and 3-bank PCI Express expansion connector.

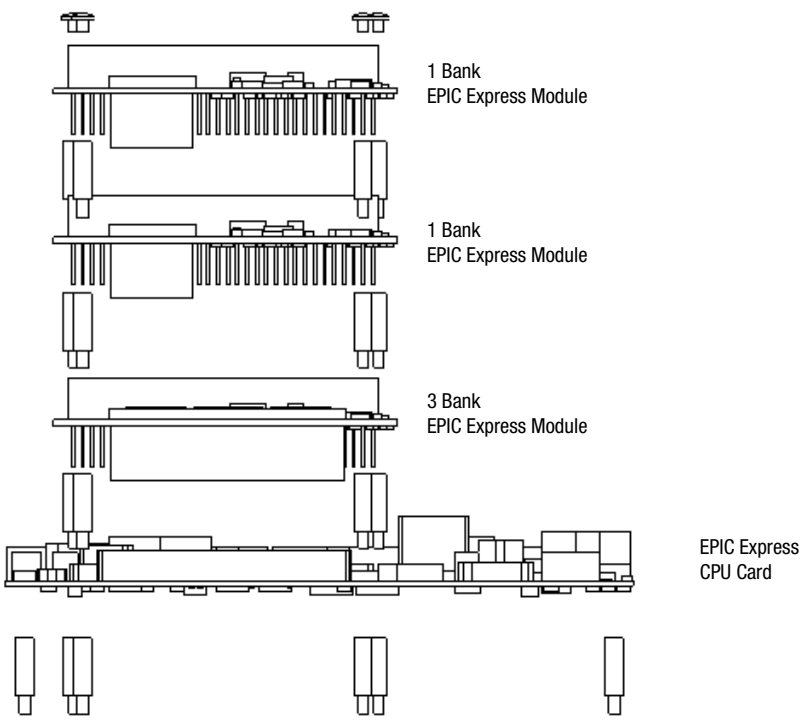

Figure 2 PC/104-size I/O modules are stackable above an EPIC Express SBC.

EPIC Express Boosts I/O Bandwidth

At merely 4.528-in. x 6.496-in., EPIC Express boards pack a lot of bandwidth punch into a tiny area. Embedded applications generally do not have the luxury of lots of space. EPIC Express defines a PC/104-size expansion interface with four x1 lanes and two x4 lanes (Figure 1).

The specification replaces the PCI/104 connector, a 32-bit/33 MHz PCI bus, with a smaller, high-density surface mount connector and 10 to 30 times the I/O bandwidth. The I/O bandwidth limitations of PC/104-Plus—1 Gbit/s on PCI—are solved.

Two plug-compatible connector options are provided by EPIC Express. The one-bank connector brings four x1 lanes, with up to 10 times the 32/33 parallel PCI bus bandwidth in each direction. The latest chipset generation on the market includes four x1 lanes, so the one-bank connector is a good choice for new SBCs and I/O modules.

The three-bank connector defines two additional lanes in the next two connector banks. Each lane is a x4 type, with four times the bandwidth of a single x1 lane. A one-bank connector on an I/O module works when plugged into a three-bank connector on an SBC.

Each x4 lane brings another 10 times the 32/33 PCI bandwidth. In total, the six lanes in a three-bank configuration are equivalent to 12 x1 lanes, or 30 times the 32/33 PCI bandwidth. The PC/104 I/O market now has an I/O superhighway (Figure 2).

PCI Express is not a traditional multidrop bus, where multiple devices wait for their turn to use the bus. The differential pair signaling cannot be cascaded from device to device. Each lane has an independent path to the chipset. This poses a challenge to a stacking architecture like PC/104.

The EPIC Express creators solved this problem with a clever totem-pole scheme. Each I/O module in the PC/104 stack consumes the PCI Express lanes that it needs, whether one lane or multiple lanes (Figure 3). The module then shifts over, or normalizes, the remaining lanes so that the next module up the stack has the lanes it needs on the connector pins that it expects (Figure 4).

EPIC Express Preserves Legacy

Since EPIC Express was created with legacy systems in mind, the hundreds of existing off-the-shelf PC/104 I/O modules do not have to be re-designed. Neither do the thousands of embedded systems deployed with PC/104 technology.

The definition of EPIC Express includes the familiar PC/104 connector and bus (ISA), so that SBCs that include the ISA bus have a standard way of bringing it out. Off-the-shelf PC/104 modules simply plug right into the SBC. For EPIC Express SBCs that do not include the PC/104 connector, a bridge card can bring PCI Express to PCI, then from PCI to ISA, so that PC/104 modules can be used the rest of the way up the stack. There is no change to the mounting hardware or I/O cabling of expansion modules within the embedded systems.

A Boost for Standards and Applications

EPIC Express is a huge leap forward in the PC/104 space. It unifies the off-the-shelf stackable embedded PC market with a single architecture. In addition, this is accomplished without sacrificing the important standards that created this market: PC/104 (ISA bus), PC/104-Plus (PCI and ISA bus) and PCI-104 (PCI bus).

PC/104 technology, including related EBX and EPIC SBC form-factor standards, is well established in medical, military, transportation, industrial, instrumentation and commercial markets. EPIC Express brings a substantial amount of I/O bandwidth to these applications without sacrificing legacy circuitry. EPIC Express also paves the way for a new class of applications that can use stackable PC/104-size I/O modules.

With 10 to 30 times the data throughput, EPIC Express boards and stackable modules are well poised for emerging embedded applications that need Gigabit Ethernet, RAID, 802.11g, custom encryption circuitry in the form of an FPGA, I/O controllers or custom high-speed circuitry.

An FPGA on an EPIC Express module can be viewed as an offload engine that frees up the CPU for other tasks. The SBC can shuttle data quickly over PCI Express lanes, complex algorithms and equations can be synthesized and programmed into the FPGA logic, and results can be returned to the SBC. High-quality

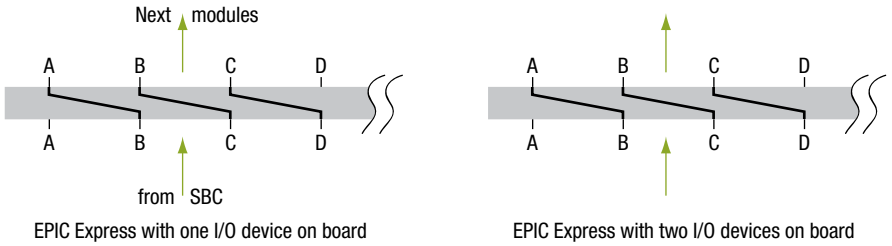

Figure 3 Unused PCI Express lanes are shifted over by the module's PCB to be in position for the next module in the stack.

EMBEDDED PLANET™

design. The next generation of connected devices.

develop. Your products based on our platform.

deploy. Your solution faster.

design. develop. deploy.

Are You Board With Me?

Start with an AMCC PowerPC 405 Processor.

Specify your I/O requirements.

Choose your operating system.

Relax while Embedded Planet delivers your rugged and reliable system.

4760 Richmond Rd / Cleveland, OH 44128
Tel: 216.245.4180 / Fax: 216.292.0561
www.embeddedplanet.com

EMBEDDEDPLANET™

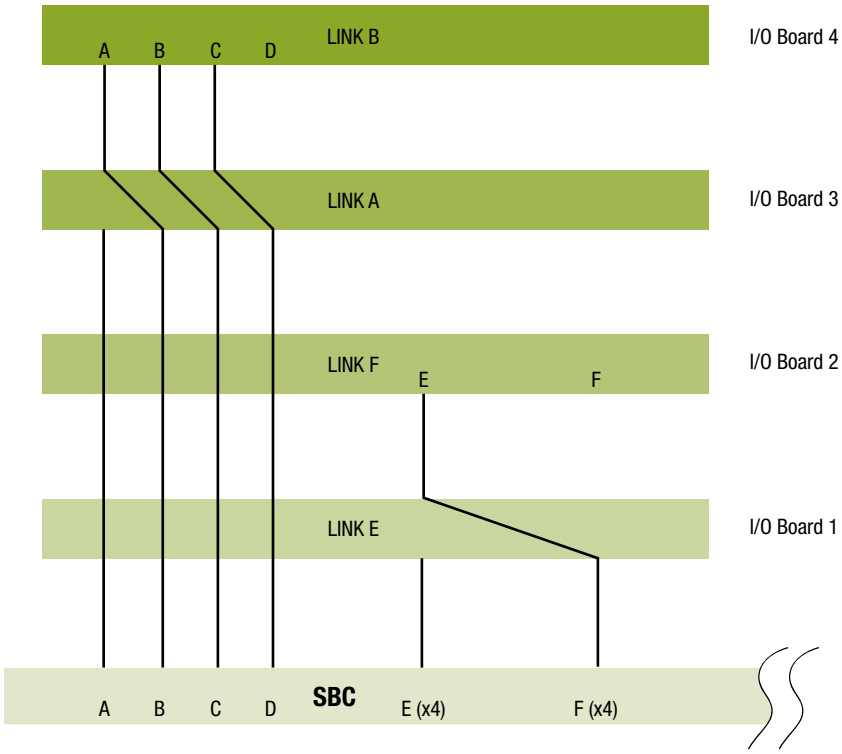

4 Board Full EPIC Express Stack with two x 4 links and two x 1 links featuring automatic alignment for each.

Figure 4 An entire stack of modules is shown, along with the PCI Express lanes used by each module.

image recognition for homeland security, both still images and full-motion video, is one example of an application that can be implemented using EPIC Express with a custom I/O module, rather than a full-custom design.

Systems that need to combine high-speed circuits with low-speed designs, such as data acquisition (A/D) and multiple serial ports, can now leverage off-the-shelf technology to improve time-to-market.

EPIC Express breathes new life into the stackable PC/104 architecture by boosting application performance while leveraging existing investments in system mechanical designs for the long haul. With a solution for point-to-point differential pair signaling, EPIC Express is next in the long tradition of tailoring motherboard technology for use in off-the-shelf SBCs, I/O modules and space-constrained systems. The first SBCs and modules are expected from suppliers in 2006. ◢

Ampro Computers
San Jose, CA.
(408) 360-0200.
[www.ampro.com].

New! Multiprocessor, Hyper-Threading and APIC support. See us at ESC 2006, Booth #1744

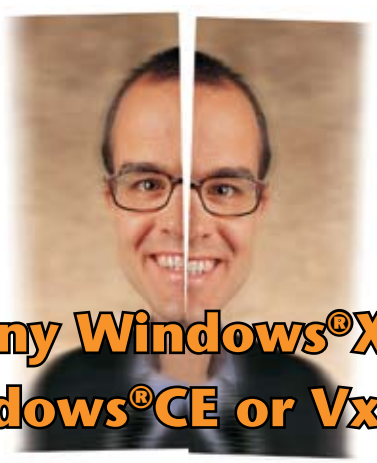

"I need hard real-time performance at the lowest possible cost."

"My customer wants high-end graphics, connectivity and a familiar interface!"

Why not run any Windows®XP application with Windows®CE or VxWorks®?

It's easy! Deliver the benefits of hard real-time control plus high-end graphics, connectivity and access to all Windows applications on your next system. Our nonproprietary, real-time extensions – CeWin® and VxWin® – enable Windows® CE or VxWorks® to co-exist with Windows XP on a single computer. The hard real-time capabilities of the RTOS are not affected by Windows XP running simultaneously. And programming is easy using well-known, proven tools. Additional control hardware or costly intelligent co-processor boards for visualization are unnecessary.

Download the free demo software proven in over 70,000 systems at... www.kuka-controls.com

The Americas • Asia Pacific • Europe

© 2006 KUKA Controls GmbH. All rights reserved. Trademarks are property of the respective owners.

KUKA Controls GmbH
Tel: 714-505-1485 • Fax 714-505-1149
Tel: 248-844-0569 • Fax: 248-844 -0505
Tel: +49 751 56122-0 • Fax: +49 751 56122 22
Sales@kuka-controls.com

KUKA

EBX

PC/104

EPIC

STD 32

Your OEM product can become extinct long before its time, unless you have a long-term commitment from your SBC vendor. VersaLogic guarantees product availability for at least five years from its introduction date.† That saves you from tossing future sales into the jaws of competitors because you suddenly cannot get your hands on the boards you need, when you need them.

VersaLogic's Extended Life-Cycle Policy is just one of the services that helped us earn the prestigious VDC Platinum Vendor award. From our close ties with key vendors to our pre-design component studies, we make it our business to extend your product's life. So sink your teeth into the details at **www.VersaLogic.com/pv13**. Start a long, long-term relationship by calling us at **(800) 824-3163.**

Before Kristin Koons joined VersaLogic's sales support team, she guided tours at a natural history museum. She really knows how to keep products off the endangered species list.

VDC Platinum Vendor

Platinum Vendor award by Venture Development Corporation 2002 and 2003

VERSALOGIC CORPORATION

(800) 824-3163
www.VersaLogic.com/pv13

©2004 VersaLogic Corp. All rights reserved.
†Complete details at VersaLogic.com/lta

DSPs, Microprocessors and FPGAs in Control

DSPs, general-purpose microprocessors, microcontrollers and FPGAs are all used in control, sometimes in combination. Each device type is best suited for a specific role, and a variety of tools is available for developing applications.

by Shelley Gretlein, Gerardo Garcia and Joel Sumner
National Instruments

As embedded technology proliferates in electronic devices used for control, it is becoming increasingly difficult for engineers to decide which technology they should select for their next designs. The silicon candidates include DSPs, general-purpose microprocessors or microcontrollers, and FPGAs. In control applications these devices are best suited for different functions. In addition, new hybrid silicon is becoming available that blurs the lines between these device types.

In general, a closed-loop control application consists of three parts: measurement by a sensor connected to the process, decision in a controller element and action through an output device or actuator.

Since many control applications require rapid processing and feedback, engineers often choose DSPs for their speed. Control applications can range from airplane landing gear devices to automotive collision avoidance systems. Studying a collision avoidance system reveals the benefits of using DSPs in conjunction with microprocessors.

Get Connected
with companies mentioned in this article.
www.rtcmagazine.com/getconnected

In this type of system, DSPs and microprocessors have distinct roles (Figure 1). The system's DSP receives input signals for rapid processing of a single task and performs tasks such as reflective phase analysis. The microprocessor performs multiple housekeeping tasks, interfacing with the host and individual peripherals.

DSPs in Control

To address typical control application requirements, a DSP contains four core components. These are data memory, program memory, the computing engine and I/O. This architecture is similar to that of other devices, including microprocessors. However, DSPs are differentiated by several additional characteristics. These include multiple, parallel pipelined datapaths; optimizations for specific vertical processing operations; buffered I/O ports; abundant registers for temporary data storage; efficient address generation for array handling; and low-power options.

These characteristics have traditionally made DSPs the silicon of choice in audio, video and communications applications for concentrated numerical processing.

DSP Programming Tools

The mathematical theory of underlying digital signal processing techniques can be fairly complex. However, the numerical operations required to implement the techniques are straightforward, consisting mainly of operations that can be performed on a command-line calculator (Start >> Run >> Calc).

The architecture of a DSP chip (Figure 2) is designed to carry out those simple calculations incredibly fast. The difficult part is understanding the architecture and implementing programs to take advantage of it. For example, to multiply and add a register/value on traditional fixed-point devices using traditional tools, engineers must implement low-level instruction bit codes and understand the mnemonic MAC, APAC and MPYK commands. However, access to the DSP registers can also be made via higher-than-assembly-level optimizing C compilers offered by DSP vendors.

Although many engineers and scientists seek higher-level tools to make the technology more accessible, programming DSP processors is usually done in C or assembler. Writing programs in high-

level tools requires less effort, but the efficiency achieved can be as low as half that of a program written in assembler.

A range of DSP tools is available to suit the skillset and application needs of most engineers, from DSP-specific vendor toolchains—usually in the form of a C compiler and assembly interface—to code generators and high-level graphical tools. Vendor tools are targeted at the traditional embedded programmer or the career design engineer who can tackle challenging, low-level designs. High-level graphical programming tools are designed for the domain expert, the engineer or scientist who wants an embedded application without learning the details of low-level implementations. These high-level tools target more general-purpose devices, such as 32-bit floating-point DSPs.

A FIR filter implemented in assembler (Figure 3), C (Figure 4) and a graphical, or "G", programming language (Figure 5) can be used to compare the tools used to program a DSP. Although specific processor architectures differ slightly, in general, assembler delivers direct access to the processor registers, C encapsulates operations into functions and routines and "G" offers more high-level abstraction and intuitive programming.

The silicon choices engineers make will depend on their skill levels. Because of the complexity of programming a DSP properly—optimized for the application and utilizing the performance available—they are not typically used for the control portion of a control application.

DSP silicon and tools are specifically designed to crunch sets of numbers, or process signals, as fast as possible. But in control applications it is more important to look at loop rate, that is, how quickly and accurately a signal can be read, a decision made and a control value passed back out. The basic needs of control applications rely heavily on I/O and accessing peripherals on the silicon. These are not the traditional strengths of DSPs.

Since the lines between some general-purpose control devices and DSPs are beginning to blur, the decision is sometimes one of personal preference and the intu-

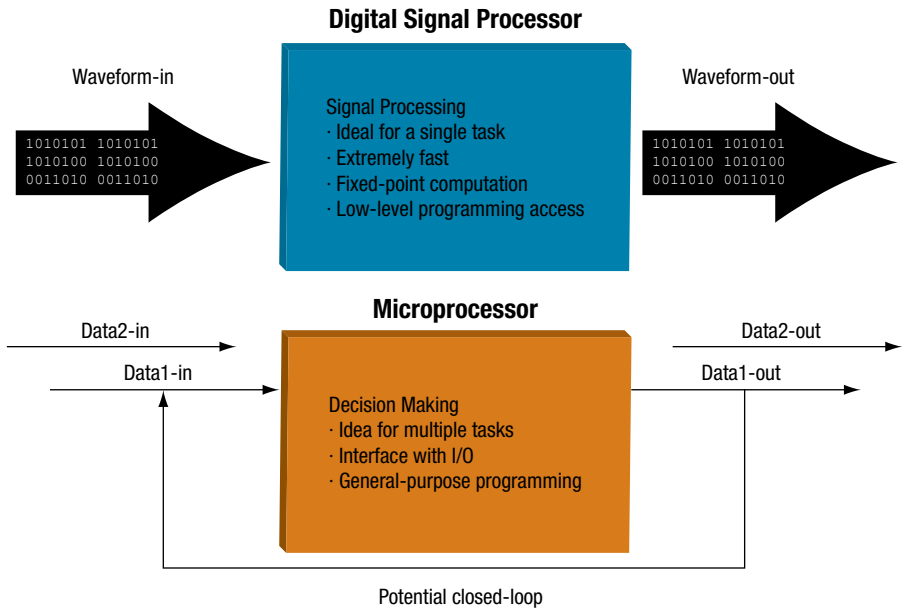

Figure 1 In control applications, DSPs and general-purpose microprocessors traditionally have distinct roles. The DSP receives input signals and performs tasks such as reflective phase analysis. The microprocessor and memory take care of housekeeping tasks for communication with the host and interface with other peripherals.

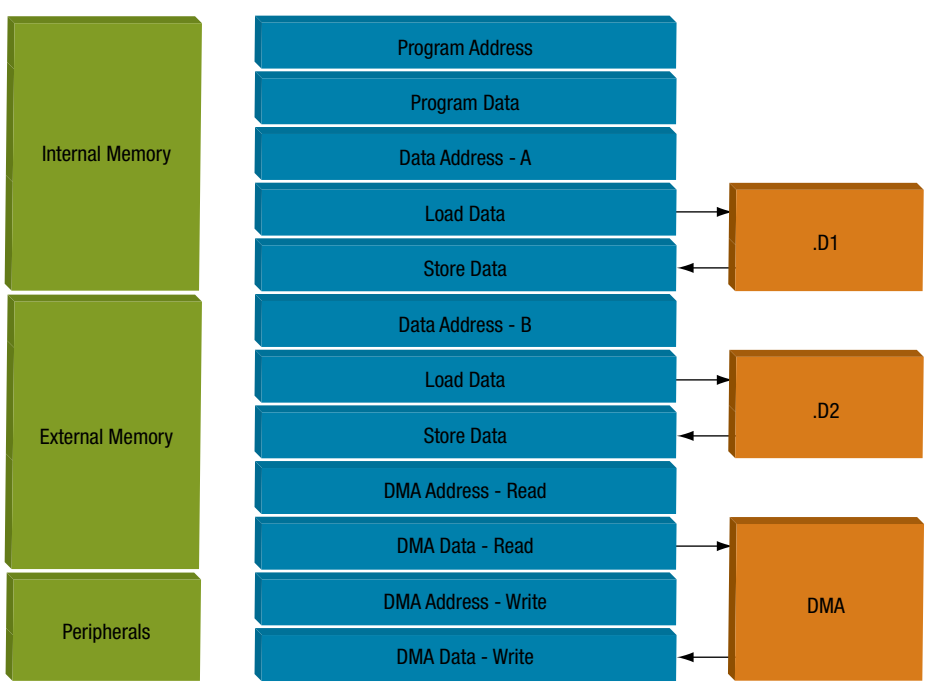

Figure 2 Many control applications require rapid processing and feedback for control. The architecture of a DSP chip, including its internal bus, is designed to conduct calculations extremely quickly, processing hundreds of millions of samples per second.

itiveness of the tool rather than one about silicon architecture.

Microprocessors and FPGAs in Control

General-purpose microprocessors and FPGAs are also widely used for signal processing and control. Microprocessors are designed to handle multiple asynchronous tasks efficiently. This is achieved by using flexible, context-saving methods and program stacks for each process that facilitate switching back and forth between tasks. Because many real-time control applications perform multiple tasks asynchronously, microprocessors are a popular choice.

General-purpose processors, as well as newer hybrid processors that combine different processor types, are the silicon versions of a Swiss Army knife. For example, Analog Devices' Blackfin processors provide both microcontroller (MCU) and DSP functionality in a unified architecture for partitioning the needs of control and signal processing. The Blackfin processor can act as a 100% MCU with code density on par with industry standards, or as a 100% DSP, or as a combination of the two. It can be programmed in assembler for the career design engineers and programmers, C for the traditional programmers and even in a graphical programming language for the domain expert.

Because they run the logic in hardware and are reprogrammable, FPGAs are gaining popularity for their speed and flexibility. An FPGA contains a matrix of reconfigurable gate array logic circuitry. When this device is programmed, the gates are configured appropriately and connected to each other in a manner that achieves a hardware implementation of the software application. FPGA-based control systems can enforce critical interlock logic and be designed to prevent I/O forcing by an operator.

Although most FPGAs used today are programmed during system manufacturing to perform a specific task, it is increasingly common for engineers to program them in the field. Unlike hard-wired printed circuit board (PCB) designs, which have fixed hardware resources, an FPGA-based control system's internal circuitry can be rewired

```
[FFA006AA] R3.L = 1164 ;
[FFA006AE] R3.H = -128 ;
[FFA006B2] P1.L = 0x774 ;
[FFA006B6] P1.H = 0xff80 ;
[FFA006BA] [ P1 + 0x0 ] = R3 ;
[FFA006BC] R2.L = 1364 ;
[FFA006C0] R2.H = -128 ;
[FFA006C4] [ P1 + 0x4 ] = R2 ;
[FFA006C6] [ P1 + 0x8 ] = R2 ;
[FFA006C8] R1 = 8 ;
[FFA006CA] [ P1 + 0xc ] = R1 ;
[FFA006CC] R0 = 1 ;
[FFA006CE] [ P1 + 0x10 ] = R0 ;
[FFA006D0] R7 = P1 ;
[FFA006D2] [ SP + 0xc ] = R7 ;
[FFA006D4] R2 = 256 ( X ) ;
[FFA006D8] R1.L = 1380 ;
[FFA006DC] R1.H = -128 ;
[FFA006E0] R0.L = 652 ;
[FFA006E4] R0.H = -128 ;
[FFA006E8] CALL __fir_fr16 ;
[FFA006EC] R0 = 0 ;
[FFA006EE] JUMP.S 2 /*0xFFA006F0*/ ;
[FFA006F0] R7 = [ SP + 0x10 ] ;
[FFA006F2] P0 = [ FP + 0x4 ] ;
[FFA006F4] UNLINK ;
[FFA006F8] JUMP ( P0 ) ;
[FFA006FA] NOP ;
           __Exit
[FFA006FC] R0 = 0 ;
           __lib_prog_term
[FFA006FE] SSYNC ;
[FFA00700] IDLE ;
```

Figure 3 A FIR filter implementation written in assembler, which delivers direct access to the processor registers.

Shatter
Processing Limits in Your Wireless Applications
Breakout performance from BittWare's New DSP Advanced Mezzanine Card (AdvancedMC™)
Combining the best-in-class performance of Analog Devices TigerSHARC with the benefits of AdvancedMC, the B2-AMC is ready for your wireless applications: 2G, 2.5G, 3G, WiMAX, and Software Defined Radio. Using the superscalar architecture of the ADSP-TS201 TigerSHARC DSPs, the B2-AMC provides an unprecedented 14.4 GFLOPs and 57.5 GOPS of processing power. A large Altera Stratix II FPGA implements BittWare's ATLANTiS architecture, seamlessly integrating the DSP processing power with your switch fabric of choice: Serial RapidIO™, PCI Express™, Advanced Switching Interconnect, GigE, or XAUI™ (10 GigE). In addition to the multiprocessor scalability of the TigerSHARC, the board features a variety of front and back panel I/O interfaces, an IPMI system management interface, and system synchronization via AMC system clocks. Ideal for use in AdvancedTCA®, MicroTCA, or custom systems, the B2-AMC will give your wireless application the power it needs.
Go ahead...
Shatter your limits.
BittWare
www.bittware.com
TigerSHARC DSP is a registered trademark of Analog Devices, Inc.

for reconfiguration while the system is deployed.

Additional FPGA advantages include the fact that their hardwired logic is extremely fast, much faster than software-based logic. They are easier to interface to the outside world, either through custom peripherals or via glue logic to custom coprocessors. They are also better suited for bit-level operations than a microprocessor.

FPGAs are especially suited to control applications requiring custom hardware, such as digital communication protocols, rapid control prototyping, hardware-in-the-loop simulation, in-vehicle data acquisition, machine control and machine condition monitoring.

In addition, different types of custom I/O hardware that benefit from the use of FPGAs include hardware for timing and triggering routines; for mixing analog I/O processing, digital I/O processing, counters and triggers on a single device; for high-channel counters, event detection, and time stamping; and for custom signal conditioning and parallel processing.

FPGAs are usually programmed with a bitstream, which configures the array of logic cells to perform the desired function and connect multiple logic cells as determined by the design. FPGA programming technologies range from one-time-programmable elements to electrically erasable devices. However, the language of FPGAs, HDL, is fairly complex. HDL is a high-level language, yet writing an HDL design description is a relatively complicated process, since it describes circuit functionality at the register transfer level (RTL).

Some engineers prefer high-level graphical tools for accessing FPGAs. For example, the National Instruments LabVIEW FPGA Module (Figure 6) lets non-embedded experts create custom I/O measurement and control hardware without the need for low-level hardware description languages or hardware board-level design. This custom hardware can be used for unique timing and triggering routines, ultra-high-speed control and interfacing to digital protocols.

A Comparison of Devices in Control

The essential difference between a DSP and a microprocessor or microcontroller is that the DSP's features are designed to support high-performance, repetitive, numerically intensive tasks. In contrast, general-purpose processors are not specialized for a specific kind of application, and microcontrollers are designed for control-oriented applications.

Engineers often choose microprocessors instead of DSPs for control applications because they traditionally offer better on-chip peripherals, such as pulse width modulation (PWM) and other integrated I/O for multiple, asynchronous tasks. Because they are general purpose, microprocessors are more flexible and less specialized, which usually makes them easier to understand and program.

Unlike processors, FPGAs use dedicated hardware for processing logic and do not have an operating system. Since FPGAs are truly parallel in nature, different processing operations do not have to compete for the same resources. As a result, the performance

Figure 4 A FIR filter implementation written in C, which encapsulates operations into functions and routines.

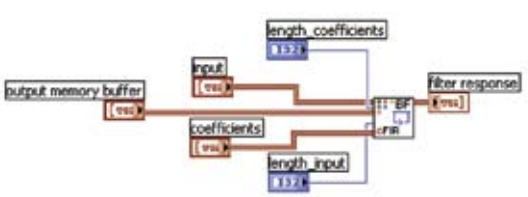

Figure 5 A FIR filter implementation written in a graphical programming language, which is more high-level and intuitive.

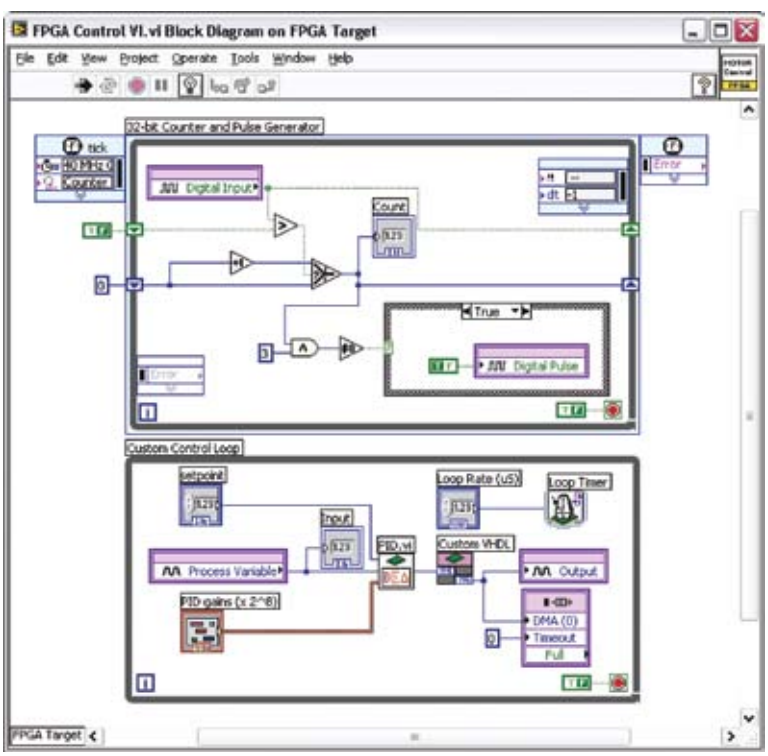

Figure 6 This NI LabVIEW graphical code simultaneously executes a 32-bit counter and pulse generator and a custom control algorithm on an FPGA.

of one part of the application is not affected by adding another parallel processing part.

FPGAs and processors also differ in compilation. When a control application for an FPGA device is compiled, the result is a highly optimized silicon implementation that provides parallel processing with the performance and reliability benefits of dedicated hardware circuitry. Because there is no OS on the FPGA, the code is implemented in a way that ensures maximum performance and reliability.

Finally, in addition to offering high reliability, FPGAs can perform deterministic closed-loop control at extremely fast loop rates. In most FPGA-based control applications, speed is limited by the sensors, actuators and I/O modules rather than by the FPGA's processing performance. For example, the proportional-integral-derivative (PID) control algorithm, commonly used for regulating analog processes such as temperature and pressure, executes in just 300 nanoseconds (3.33 MHz), whereas a PID loop executes at the rate of 2.75 kHz on a typical industrial processor.

However, FPGAs by themselves are not as useful as when they complement microprocessors or DSPs. Interfacing an FPGA to the outside world through A/D controllers, D/A controllers and a microcontroller offers complete system designs with fast I/O, appropriate control and rapid signal processing.

Choosing the Right Device for the Application

Although these differences distinguish DSPs, microprocessors, MCUs and FPGAs from each other, as with most decisions, an analysis of the particular application must be done to discover which device or combination of devices best suits it.

If the application must access a variety of I/O, reprogram logic or run closed-loop control at rates exceeding 1 MHz, an FPGA is probably the right fit. If designers want general-purpose functionality and programming simplicity, a microcontroller is the more appropriate choice. If ultimate optimization and processing speed are needed and the engineer is familiar with underlying chip architecture and DSP programming, a DSP is the best device to use.

Finally, if engineers do not want to learn low-level programming, the tool used to access the silicon may be the determining factor. In this case, an integrated graphical programming tool may be the right choice. ◢

National Instruments
Austin, TX.
(512) 683-0100.
[www.ni.com].

Total Rugged Solutions for Extreme Environments
Trusted ePlatform Services
ADVANTECH
Advantech Stackable SBCs that Provide Rugged Features for Total Solutions
■ Extended Temperature Testing (ETT) Services
- Phoenix Operation, for Wideband Temperature
- Phoebus Design, for Extreme Temperature (-40 ~ 85°C)
■ Conformal Coating Service
■ Glued DRAM Service
PCM-4386 EPIC Celeron® M Solutions
PCM-3380 PCI-104 Pentium® M Solutions
PCM-3341 PC/104 STPC Altas
PCM-9582 EBX Pentium® M Solutions
Advantech Corporation
15375 Barranca Parkway, Suite A-106
Irvine, CA 92618
Toll Free: 1-800-866-6008
Ph: 949-789-7178
Fax: 949-789-7179
Email: ECGInfo@advantech.com
Website: www.advantech.com
© 2006 Advantech Co., www.advantech.com

Cool. Calm. Collected.

The Themis Slice™ architecture can handle whatever you or the outside world throw at it.

Cool.
Innovative and cost-effective cooling system can tame the hottest of processors in four 100 Watt sockets that support eight processor cores. Air and liquid heat exchanger options.

Calm.
Standardized mechanical footprint designed to survive in high shock and vibration environments - up to 40G shock and 1 grms vibration.

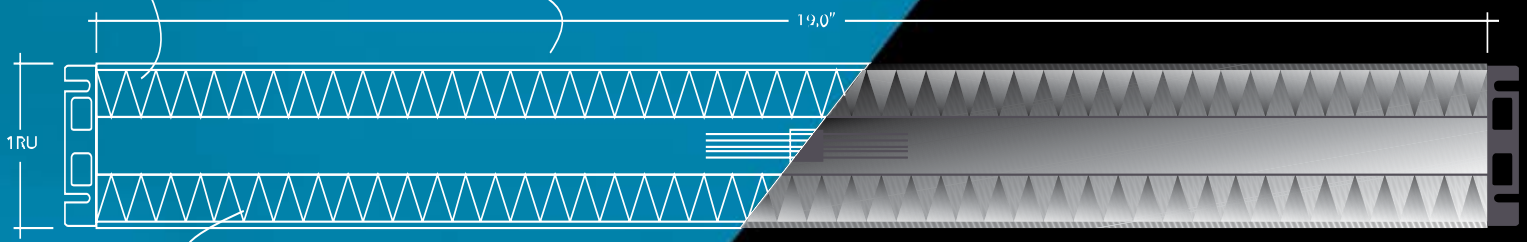

Collected.
Scalable - up to 16 processor cores. High-density, processor-independent architecture implemented with SPARC®, IBM®, PowerPC and x86 technology. Solaris™, Windows®, and Linux® operating systems. All networked in a single system. All managed in real-time.

Themis' Slice technology is a breakthrough in lifecycle management for mission-critical computing.

By functionally disaggregating commercial computing resources and housing them in a standardized footprint, purpose-built enclosure, Slice provides resilience with superior thermal and kinetic management. This open and modular design allows for spiral technology refresh, extending computing infrastructure investments for complete lifecycle management.

Themis' real-time resource manager integrates all computing resources, under dynamic policy control, to optimize application quality of service.

Themis Slice is transformational technology offering the most robust and efficient platform available. No other platform will enable the 21st century war fighter as effectively.

www.themis.com (510) 252-0870

THEMIS

Transformational.

© 2005. Themis Computer, Themis, Themis logo and Themis Slice are trademarks or registered trademarks of Themis Computer. All other trademarks are property of their respective owners.

DSPs + FPGAs Provide Flow Cytometry and Cell Sorting Solution

Flow cytometry and cell sorting systems have traditionally utilized complex, custom analog circuitry. An intelligent system based on DSPs and FPGAs, combined with the right architecture, delivers enough information to reduce false and missed detections, as well as a high level of flexibility and reprogrammability.

by Kristen K. Zaffini
BittWare

Until recently, many inspection and control systems have required a human in the loop to provide the actual inspection. However, due to the inefficiency, and sometimes ineffectiveness, of manual inspection, automated inspection and control has become the most sought-after method. This has proven to be no easy task, especially in the complex world of flow cytometry and cell sorting. This is the process used to measure the physical or chemical characteristics of a biological cell and, using this information, sort cells from a sample. Such high-end applications require an intelligent system that can also provide a high level of flexibility and reprogrammability, so that the system is, to some extent, a general-purpose one that can be tailored to the needs of individual applications.

For years, DSPs have been the processor of choice for applications requiring the type of compute-intensive, high-speed calculations used in both medical and bioscience fields, as well as in manufacturing. FPGAs are now being touted as do-everything processors, especially when reconfigurability is a requirement. Often, the best solution is a combination of FPGAs and DSPs working in conjunction with each other (Figure 1).

Basic Inspection and Control System Needs

All systems used for inspection and control, regardless of industry, include two stages of processing. The first stage is the initial inspection during which the system acknowledges the presence of a sample. In the second stage, the required feature extraction is performed. This consists of gathering certain characteristics from the unit under test, often using digitized images for pattern recognition and feature extraction; providing the "intelligent" processing of these images; and ultimately determining if this is a "good" sample or a "bad" sample.

The processing required for the initial inspection, although fairly minimal, needs to be repeated indefinitely as the system consistently monitors for the next unit. Regardless of the type of control and inspection system, the second stage of processing is quite complex.

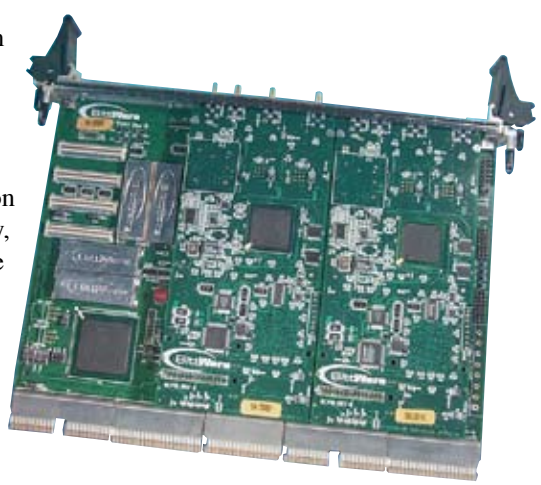

Figure 1 Combining BittWare's 6U CompactPCI DSP board with two of its high-speed Barracuda-PMC+ data acquisition cards provides eight TigerSHARC DSPs and three FPGAs. The combination of FPGAs and DSPs working in conjunction with each other is well suited to applications that require different levels of processing, such as a flow cytometry and cell sorting system.

Get Connected
with companies mentioned in this article.
www.rtcmagazine.com/getconnected

Flow Cytometry and Cell Sorting Systems

In high-end bioscience applications, automated inspection and control is often required to perform tasks that would otherwise be impossible, since they cannot be done manually. The sorting and classification of biological cells via flow cytometry and cell sorting is one such area. This type of control and inspection system is very similar to the basic system described above in that there are two stages: the initial constant inspection to determine the presence of a cell, and the intelligent processing. The end goal of the system is to divert a droplet containing a cell into the appropriate receptacle by analyzing the cell's identifying characteristics.

Before each droplet is released from the sample, the system must identify that a cell is present and begin to analyze it. In order to analyze and isolate certain characteristics of a cell, a fluorescent compound is formulated specifically to highlight those characteristics. When added to the sample, this compound binds itself to the cells that possess the desired characteristic. The pressurized cell suspension is forced into a tube and then through a needle, making the cells form a single-file line. The line of cells is then forced through a small nozzle, emerging as a jetting stream. The jetting cells pass one at a time through a laser beam where the cells scatter laser light.

If the desired characteristic is present, the compound in the cell also emits fluorescent light. This fluorescence is detected by a photo sensor, which converts the fluorescent light into a momentary voltage pulse. This stream of voltage pulses is fed into an A/D converter that continuously samples the stream of pulses. Within microseconds, the system needs to perform the pulse detect, identifying that a cell is present, then extract the desired features from the pulse, analyze these features and ultimately determine whether or not the cell is to be sorted.

These tasks take place in parallel so that the system may, at any single point in time, simultaneously detect one cell, analyze another cell, and route yet another cell. This requires a microsecond-level timeframe. Within one second, there could be up to 100,000 cells shooting out of the system, all needing to be measured, classified and routed. Typically, the sorting decision must be made 100 microseconds after the cell passes through the laser beam.

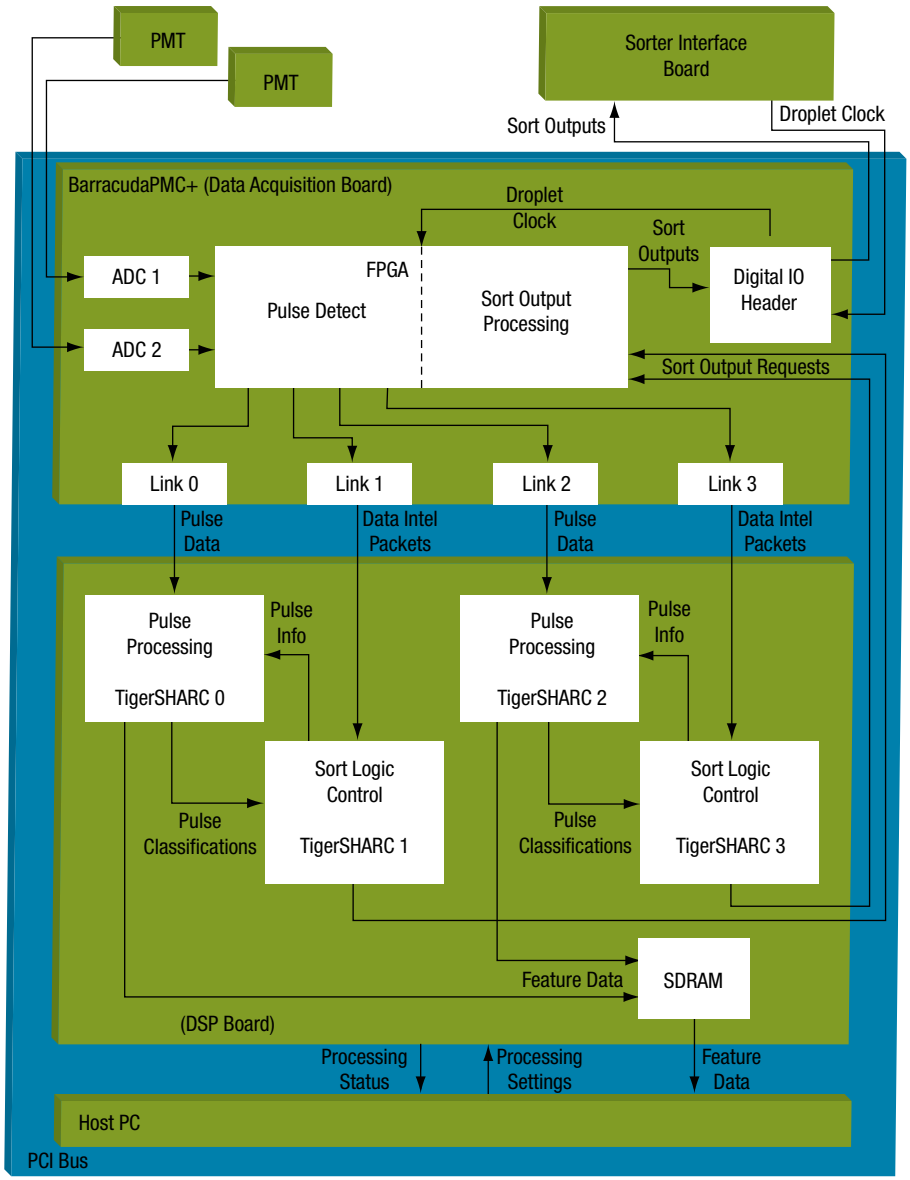

Figure 2 A flow cytometry and cell sorting system measures the physical or chemical characteristics of a biological cell and uses that information to sort cells from a sample. This design uses two high-speed A/D cards, each including an FPGA for the initial stage of processing, and one DSP board for the high-end processing required by the second stage.

Flexible, Intelligent Cell Sorting

These flow cytometry and cell sorting systems have traditionally utilized complex, custom analog circuitry that has remained unchanged for many years. One manufacturer, iCyt Visionary Bioscience, decided that it was time for a change. The company wanted to design an intelligent system that would also provide a high level of flexibility and reprogrammability, enabling it to be somewhat general purpose, yet capable of being tailored to the requirements of specific applications.

The older systems sampled only the maximum peak fluorescence coming from a particular cell, providing merely a single

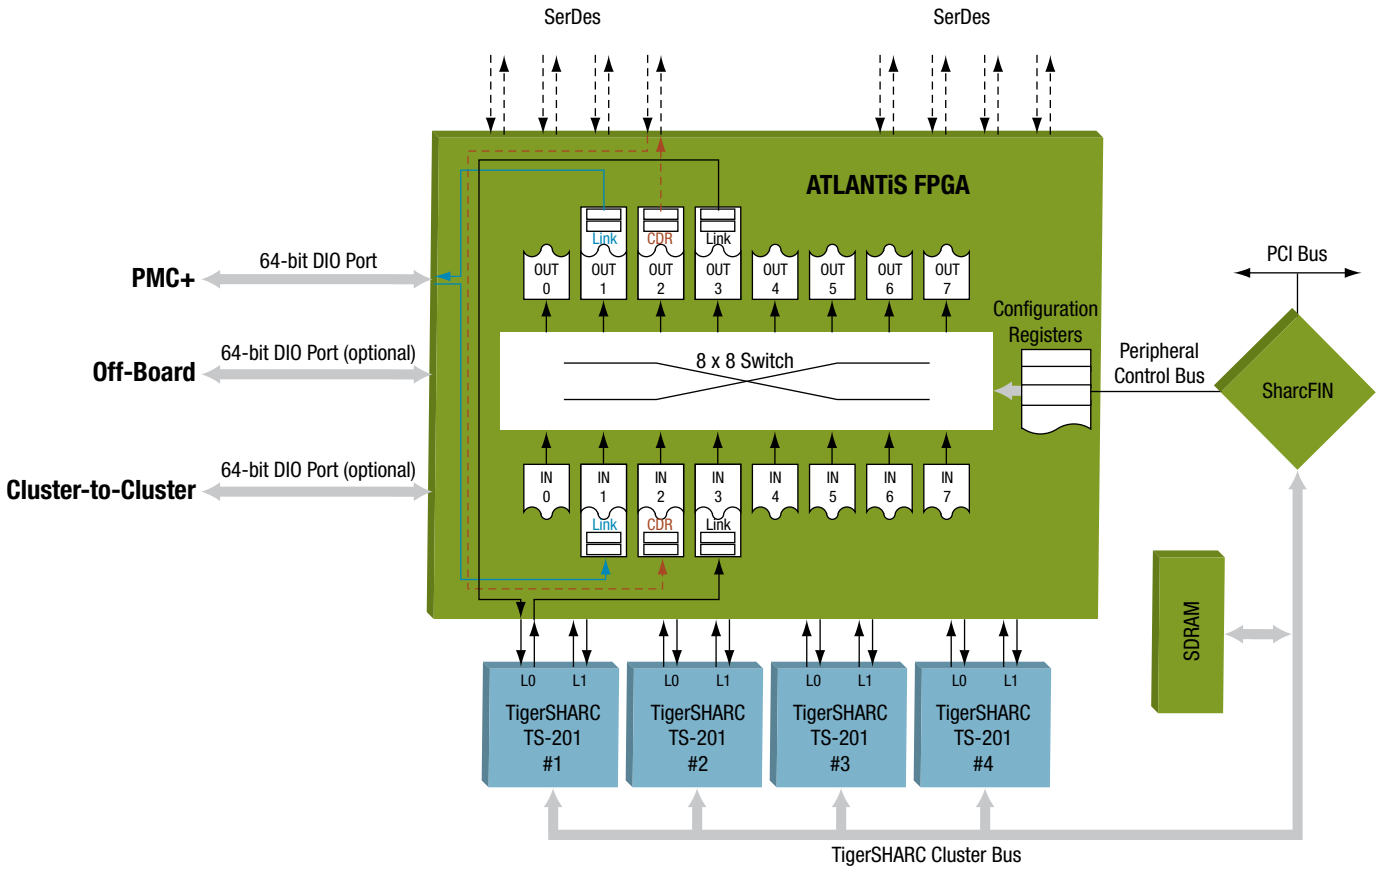

Figure 3 The ATLANTiS architecture is implemented within each onboard FPGA. It comprises one or more switches connected via transmit and receive blocks to the onboard I/O. The configuration registers are used to program the interconnectivity between all inputs and outputs and to add processing blocks if needed.

data point. The proposed new design would capture at least 100 samples per cell, giving the system more information to analyze, such as rise times, fall times, width and integrated energy. This additional data would greatly reduce false detections and missed detections.

A cell typically passes through the laser beam in about one microsecond. Within no more than 100 microseconds, the droplet containing the cell must be routed to the correct receptacle. During this time, the system must detect the presence of a cell, extract its relevant features, analyze this data and sort the cell. This requires a sampling rate of at least 100 MHz. At 2 bytes/sample, the required sustained throughput is 200 Mbytes/s. The technology chosen had to support this throughput while also providing extended precision and dynamic range on unpredictable and changing data in a fixed amount of time.

Although the proposed system would be targeted at cell sorting, many specialized applications are possible within that field. The system would therefore need to be flexible, enabling not only a variety of uses, but also satisfying requirements that do not yet exist.

A Choice of Processors

When determining which technologies should be utilized in the design, the company's designers had to take into account four separate requirements: low latency, high throughput, complete flexibility and the system's ability to be rapidly adapted and its design updated.

The question then became which processing technology to use, from a traditional microcontroller or microprocessor, to a DSP or programmable logic device.

The speed and dynamic range required for the system pointed to a DSP. They deliver sustained, low-latency, high-throughput data processing. The additional dynamic range provided by a floating-point DSP, required by this proposed system design, also ensures that the system will easily support future applications. A general-purpose microprocessor would be hard-pressed to match this set of abilities. Flexibility, unmatched algorithmic capabilities optimized for fast floating-point math, and the benefit of programming using C, effectively decreasing development time, are all benefits of designing with a DSP.

DSP-Only Design vs. Hybrid DSP/FPGA-Design

The prototype used four high-end, floating-point TigerSHARC DSPs per channel to accomplish the first and second stages of processing in the flow cytometry and cell sorting system. These ultra-high-performance, real-time multiprocessing DSPs were chosen for their high throughput, low latency and deterministic characteristics. They provide continuous complex processing at high speeds in a fixed amount of time, crucial to this application.

After further consideration, it was determined that the initial stage of processing could be accomplished much more efficiently with an FPGA. Utilizing a lower-cost FPGA for the pulse detect would save the interesting data for the DSPs, which are particularly well suited to the type of high-throughput, data-stream processing required in the second stage.

The FPGAs, on the other hand, could easily perform the processing required for the initial cell detect. This data does not require complex computational math, but

does require fast, repetitive processing. Since the DSPs are not customized for this type of data analysis, they were wasting large amounts of processing effort on the first stage. Although the FPGA processing only accounts for 20% of the system's total processing, the number of DSPs required was cut in half once the FPGA was added.

Significant design work would be required to ensure that the FPGA and DSPs operated efficiently together. Many design engineers are accustomed to working primarily with either DSPs or FPGAs, but not both, and might be unfamiliar with the requirements of integrating the less familiar silicon into a design. An architecture which provides that for designers becomes crucial. This is especially true if it also enables them to easily change the routing of inputs and outputs of the signal processing system for future expansion and modification.

The final subsystem (Figure 2) came in the form of BittWare's hybrid embedded design, which uses the company's custom ATLANTiS architecture to interface between the FPGA and DSPs. This architecture is an I/O routing scheme in which every I/O can be dynamically connected to any other I/O. By employing this architecture and utilizing both DSPs and FPGAs in the design, a powerful, completely flexible system was created. The FPGA provides the initial processing, performing the pulse detect, and passes the data along to the DSPs only when a cell is actually present. The DSPs then provide feature extraction and intense data analysis.

The architecture (Figure 3), implemented in the FPGA, provides one solution for the three dilemmas typically faced in multiprocessor designs. These are: how to allocate the I/O bandwidth among the processors, how to easily connect the various I/Os together while retaining the ability to modify them, and how to integrate the FPGA processing with the DSP processing. It solves all of these dilemmas by enabling communication between the TigerSHARC link ports and all other I/Os connected to the board and by providing continuous 2 Gbyte/s throughput. This communication can be point-to-point, or one input can broadcast to all or multiple outputs. The I/Os can be connected or disconnected from each other as requirements dictate without the need for recompiling or changing cables.

This architecture consists of one or more switches connected to a configuration register that is controlled by the designer using a GUI. The routing can be changed at any point by reprogramming the configuration registers. Since all I/Os connected to the board are input into the FPGA, and thus into the architecture itself, standard and/or custom FPGA processing blocks can be inserted in the queue at any point during the transfer of data. For both novice and seasoned FPGA designers, this architecture enables seamless data transfer and the option of including processing blocks.

The final system using both DSPs and an FPGA, coupled with an architecture that gives designers the ability to easily integrate the two, provides an intelligent and flexible flow cytometry and cell sorting system. The design easily handles the initial system requirements, while also providing enough flexibility for future enhancements and upgrades.

BittWare
Concord, NH.
(603) 226-0404.
[www.bittware.com].

at the RTECC Boston
Sheraton Framingham Hotel

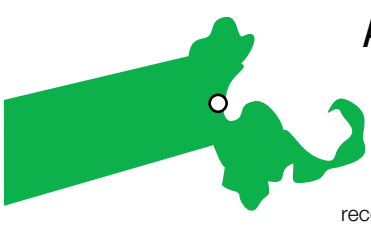

April 27, 2006

Join us at the Real-Time & Embedded Computing Conference (RTECC). Learn why, for years, DSPs have been the processor of choice for applications requiring compute-intensive, high-speed calculations -- yet now FPGAs are being touted as do-everything processors (especially when reconfigurability is a requirement). Talk to the experts from Bittware and find out how the best solution is a combination of FPGAs and DSPs working in conjunction with each other. Go to www.rtecc.com/boston to pre-register for this complimentary conference.

Real-Time & Embedded computing conference

www.rtecc.com

Windows Embedded Windows Server System

OEM Technical Seminars

presented by Microsoft

Everyday developers are choosing Microsoft® Windows® Embedded for their next great device.

See for yourself how Microsoft Windows Embedded can help you optimize device design, enhance reliability and integration, and improve time to market and development costs.

These sessions include a presentation and technical demonstration of the following Microsoft Windows Embedded Products, plus time to network at the exhibition.

- Windows CE 5.0
- Windows XP Embedded with Service Pack 2
- Windows Embedded for Point of Service
- Windows Embedded Server Products
- Special Guest Presentation by intel

Win a FREE Portable Media Center!

register today www.microsoftembedded.com

Attend this event and receive a FREE evaluation copy of Windows XP Embedded or Windows CE.

2006 Strategic Locations

Mountain View, CA January 24	Reston, VA April 11
Redmond, WA February 24	Boston, MA May 23
Chicago, IL March 21	Atlanta June 20

Exhibitors

Windows Embedded Partners

ADVANTECH Applieddata.net Ardence ICOP LOGIC

Microsoft Authorized Embedded Distributors

Arrow Electronics, Inc. AVNET Bell Micro bSQUARE

*Content, topics and exhibitors may vary by city. Please visit www.microsoftembedded.com for complete information.

an RTC Group production

Executive Interview

RTC Interviews Len Crane, President and CEO, VersaLogic

***RTC*: PC/104 was first developed sometime in the late 1980s and formally standardized in 1992, making it one of the longest surviving bus architectures (except for VME). And, it's continued to grow every year and explosively within the past five years or so. To what do you attribute the successful longevity of PC/104? Can the success continue?**

Crane: I believe that the success of the PC/104 platform is due primarily to the wide support that it has. There are additional reasons for its success, including its small size, low-cost physical configuration (self-stacking, no backplane needed), ruggedness and ease of interfacing (for custom add-on functions), but the wide variety of suppliers is the one that I would put at the top of the list. You can purchase high, medium or low-performance CPUs and marry them with an amazing array of additional I/O (Ethernet, CAN Bus, analog, motion control, etc.) right off the shelf. You can have the whole system in your hands tomorrow. That's hard to do with most board formats. And unlike more proprietary formats, OEMs ultimately have the safety net of ordering from alternate suppliers if they ever need to. If supplier X has their plant burn down and can't get back on line, customers can contact supplier Y for a similar product and continue their production with minimal disruption. In the real world that is quite valuable!

PC/104 products are also used to expand EBX and EPIC format boards, which broadens the scope of PC/104 compatibility and industry support even wider.

Of course most "PC/104" boards these days are really "PC/104-Plus" products. The evolution of the platform to use the higher-speed PCI interconnect is another reason that it has remained strong among so many other form-factors. Of course, this year we will see the first EPIC Express products released, which will mark yet another evolutionary step for the PC/104 format.

***RTC*: Referring to question one, different analysts have differing market figures for PC/104's growth and current business volume. Growth rates range from 15% on the low end to as much as 30% CAGR on the high end, and total business (including EBX, etc.) from a low of $100 million to figures topping $300 million. Can you share your evaluation of the market both in terms of growth and volume? What do you predict for the future?**

Crane: We see the PC/104 market continuing to grow, with our growth forecast around the 20% mark. I don't believe that my guess at the total market size is better then anyone else's, but it seems clear that the number of applications that need to use this kind of technology is not likely to decrease any time soon. OEMs are almost required to automate new devices that they design, and few organizations would benefit from trying to design a single board computer themselves. The only question is how many of them will be low, medium or high-volume applications. Off-the-shelf PC/104 products shine in the low and medium-volume range, but custom designed products make more sense at high volumes. What will the mix be in the future? I'm guessing that the mix will stay about the same, with a continuing increase in the total number of devices and embedded systems to run them.

***RTC*: COM Express has been getting its share of publicity recently. And, while it represents a slightly different approach to open modular systems, it is still judged by some as competitive with approaches such as PC/104 and PC/104-Plus. What is your view on how the two compare? What are the strengths and weaknesses of each approach?**

Crane: I see quite a difference between these two platforms. COM Express is not a single board computer, but a system module that must be used with a base board or "carrier card." It is great for applications where the creation of a custom

Get Connected
with companies mentioned in this article.
www.rtcmagazine.com/getconnected

carrier card is not considered a handicap (cost or time wise). This is not so great for most low and medium-volume applications. More importantly, COM Express is designed for very high performance.

Higher performance and higher speed signals mean a higher cost to the customer, and even more challenging to create an interface to the system board.

I see off-the-shelf PC/104 products continuing to fill a majority of the needs for small and medium-sized projects. It's just more cost-effective to use a ready-to-go product, unless you're looking at significant product volume. You also have to consider the development turnaround and time-to-market. And I suppose I should mention that PC/104 is a smaller size.

COM Express will be a welcome addition for those needing very high bandwidth and performance, and working on projects large enough to design them in. These would not have been PC/104 projects to begin with, so I really don't see much overlap or competition between these platforms.

***RTC*: The European Union's Restriction of Hazardous Substances (RoHS) directive (2002/95/EC) has put fear into many manufacturers and users of embedded computers both here and in the EU, perhaps more because of logistics issues than technical issues. However, despite all the claims of obstacles to compliance, VersaLogic has been one of the first to announce a RoHS-compliant PC/104 single board computer. Can you share with our readers some of the major obstacles in implementing RoHS compliance? What was your approach to part numbering, labeling, certification and testing? How did you devise systems that would be compatible with both your suppliers and your customers? And how long did the transition take?**

Crane: Wow. How much space do you have? I could write a good-sized book on our RoHS effort. I'll try to give you the short version. RoHS is a good match for VersaLogic in that environmental stewardship is part of our culture (actually, part of our mission). We endorse the concept behind RoHS and efforts to keep lead out of landfills. We embraced RoHS early on, and scheduled several new designs as RoHS-only products. What we found was that everyone (i.e., the component manufacturers) was waiting for each other, and the assembly houses were waiting for them. Until a few months ago, it was impossible to build a RoHS board because there were so few RoHS-compliant components available.

The actual roll-out of RoHS parts (when it finally came) had been chaotic in many aspects. For example, some manufacturers realized that they needed to offer RoHS components, but they couldn't manage the logistics of distributing both the existing leaded parts and a new line of RoHS parts. In fact, some manufacturers couldn't even manage the task of creating new part numbers for hundreds of new parts, so they didn't. We soon found ourselves receiving both leaded and RoHS parts under the same part number. Try to sort that one out! Another area of frustration was the lack of standardization in the manufacturers' material declaration sheets. It often took weeks to obtain a sheet, only to have it be so unclear when we got it that we had to start the processes all over again (or disqualify the part).

There were humorous moments as well. One manufacturer changed the leads on their leaded part to a RoHS-compliant metal and started shipping the new part. We built the part into some prototypes, which didn't operate at all. We traced the problem to the component in question. The manufacturer hadn't considered that RoHS products have to be assembled at higher temperatures than the leaded boards. The body of this particular component wasn't rated for RoHS temperatures, and it was simply burning up as it went through the reflow oven on the assembly line.

Ultimately, we ended up spending unimaginable numbers of staff hours doing research to figure out how to handle and assemble RoHS parts in a way that would yield identical or better quality to the leaded processes. We also addressed identification, inspection, segregation (separate inventory and repair areas), assembly, rework, etc. We made extensive changes to our production floor and our ERP system, including the ability to attach a Materials Declaration sheet to each manufacturer's part number so that we could document and certify each part.

In the end, it still comes down to whether the RoHS parts you need are available. Our first RoHS release was ready last year, but the release was delayed for four additional months waiting for a single component (a connector) when the factory had trouble producing the new RoHS version. You just can't build a board unless you can get *all* of the parts!

***RTC*: Referring to the above, there are many that claim that products based on some of the older specifications will experience difficulties as many of the ICs such as PCI interface chips go end-of-life and are no longer supported by major chip makers such as Intel. How do you propose to get around some of**

these problems as you continue to expand your RoHS-compliant family of products?

Crane: Yes, the introduction of RoHS parts has caused major upsets in the supply of leaded parts. Again, the ability of component manufacturers to simply double their product line (i.e., both leaded and RoHS product lines) was understandably limited. As RoHS parts are being released, we've seen a significant number of leaded parts that have gone EOL much earlier than planned. This is a major problem in our market segment, as longevity is often the name of the game. The most significant issue that we've seen with the RoHS rollout, is that we can no longer make many of our existing (leaded) products! The solutions to this are many, but they generally revolve around purchasing LARGE amounts of components to keep a product going, or redesigning the board. Both options have plenty of drawbacks. The most frustrating part of this is that we can't seem to get a handle on what is going to be available and what isn't. We just designed two new products, and before they made it to their Beta run, both had parts go EOL, so that we could not manufacture the board. While we are revising a product we'll have yet another part go EOL, so it's quite a futile exercise. Of course we're not yet getting a good supply of RoHS parts yet either. It appears that the only sane solution is for board-level manufacturers to take a year off, and come back when things have settled down and we know what components are really going to be available. Could you arrange for a mass vacation?

***RTC*: PC/104 and PC/104-Plus have managed to keep power budgets down while increasing performance for many generations through the use of low-power, slower processors and clever design. However, as products move to greater levels of performance and include such standards as EPIC and EPIC Express, it would seem likely that such systems would call for higher power budgets. How do you plan to deal with the higher power and thermal budgets?**

Crane: From my viewpoint, the CPU manufacturers have really been quite responsive to concerns about power consumption. Although it's usually painful to generate all the power rails required to accomplish it, the newer processors seem to be holding the line on power usage, even though performance is increasing. The Pentium IV to Pentium M is a good example. The fact that a PC/104 processor board can be sandwiched between two other PC/104 boards (resulting in little airflow to the CPU), has always be a challenge, but I don't think that thermal management issues are going to change radically in the near future.

***RTC*: While PC/104 is used in a very wide range of embedded computer applications ranging from POP retail to medical instrumentation, is there any one application—or family of applications—that stand out in your mind as the "ideal" application(s) for PC/104, PC/104-Plus, EPIC? Why?**

Crane: "Ideal" in my mind comes back to two primary factors: small size and inherent ruggedness (lack of board flex). Avionics and military applications come to mind as having space constraints and ruggedness high on the list. But you certainly see the same needs in applications from medical to transportation to instrumentation and even factory automation. There are few applications that don't benefit from the system controller being able to fit in a small space and be very rugged.

The tradeoff, at some point, is the increased cost of packaging additional functionality in a very small package. The assembly, packaging and cabling of a PC/104 system can take some effort, depending on the size and complexity of the system. For applications that have more available space, or liner space versus cubic space, the EBX or EPIC formats provide an advantage. These bolt-down formats are typically easier to install than a stack of PC/104 boards, provide about the same functionality, and still allow PC/104 expansion if needed.

***RTC*: PC/104 has recently had a strong showing in the industrial control market. Back in the early 1990s, low-cost, Asian-made PCs made a significant impact on the market virtually eliminating many of the older VME and STD products. However, there seems to have been a rebound in favor of industrial versions of embedded computers in the past several years with PC/104 family products taking up most of the slack. To what do you attribute this rebound? Need for rugged platforms? Size? Cost?**

Crane: Low-cost Asian PCs may have been placed in some applications that competed with PC/104, VME or STD, but we didn't see a lot of this with our customers. As with most markets, there are customers for many flavors of products. The low-cost products were great for some markets, while other factors were more important to others. In our market segment, product reliability, stable designs (long-term product availability), access to technical support during the OEM development phase, and other factors, have been much more important than the per-piece product price.

For example, if you are a medical manufacturer and you finally make it through clinical trials and FDA approval, the cost to start all over (because your embedded computer board is no longer available and you have to design in a new model) is quite high. If you are installing a control system in a high-vibration environment, and you have to repair the system every month or two because it shakes itself apart, the cost of the product is not going to be your main concern (as opposed to the cost of service calls). Expert, immediate technical assistance during an OEM's product development cycle can make a huge impact on the time-to-market. I think that many project managers are more aware of these factors today, and the real costs of various approaches are weighed up front. In some applications a lower-cost product will work just fine. In others they can end up being very expensive.

***RTC*: Over the past five or six years or so, there's been tremendous consolidation within the embedded computer industry. At least 20 acquisitions have occurred and it's reported that more are on the way. At least part of the ra-**

tionale for many of these acquisitions is so larger companies can control more and more of the product that goes into a system. However, to date, we've seen little of this on the PC/104 front though we've heard rumors of some pending transactions. Is there something that makes PC/104 companies immune, or at least resistant to acquisitions of this kind? Why do you think this has not been an issue with PC/104 vendors? Do you think it will be in the future?

Crane: I think that PC/104 manufacturers in general may be a little different than other embedded vendors (in a good way). They serve a fairly small specialty/niche market and are often lead by engineers, rather than investment bankers. These are mostly closely held organizations that are more interested in customer needs and product development than quarterly growth targets or cashing out. In general, they are more nimble and responsive than larger organizations could be, and they usually enjoy their work!

PC/104 companies are generally smaller, which may make them less of a target for acquisition. Or perhaps they don't believe that bigger is necessarily better and don't want to go through the pain of merging two corporate cultures into one. Whatever the reason, they seem to say "no" when buyers come knocking at the door, and we are left with a great variety of PC/104 vendors.

In my view, the wide base of support for a market like PC/104 and the healthy competition that it brings, is vital to its continued success. Consolidating many of the PC/104 companies into one or two would undoubtedly diminish the variety of products and services that are available, and probably shrink the market as a whole. I think it really comes back to PC/104 being more of a specialty market with a rather entrepreneurial style to it.

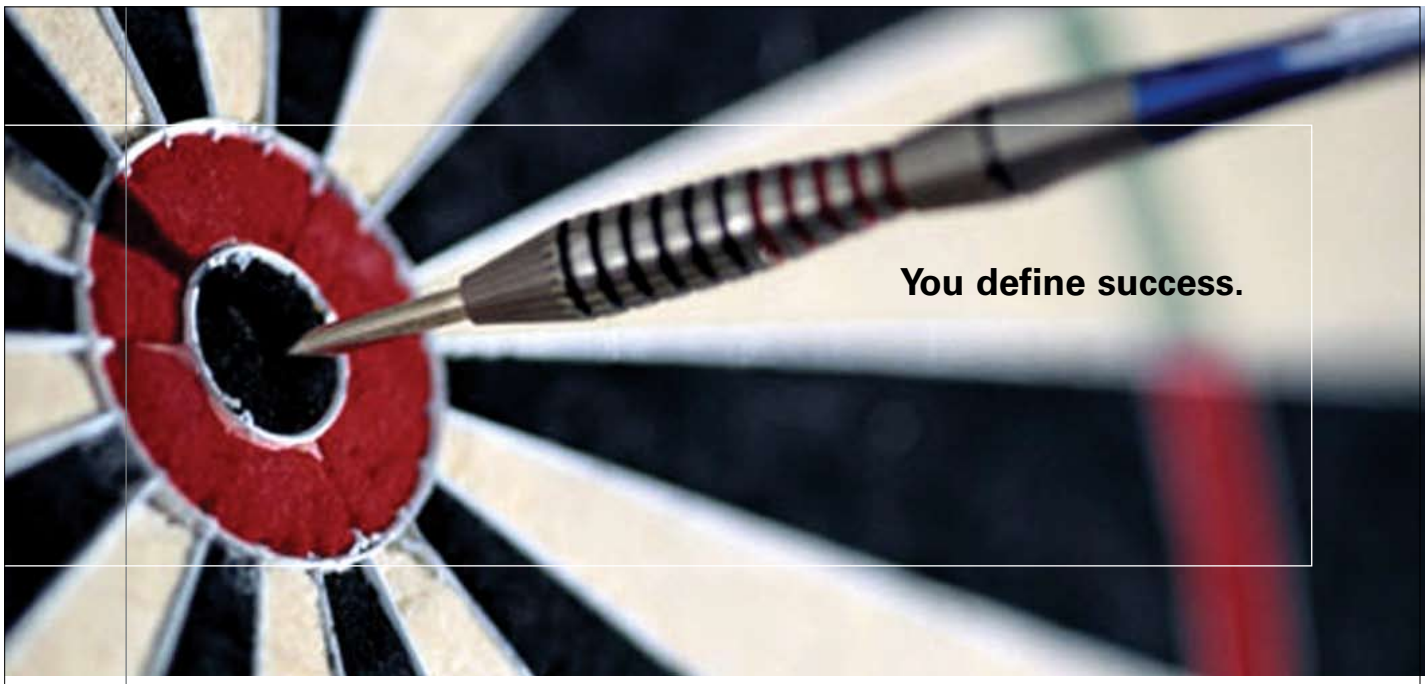

We design solutions.

Mercury customers span multiple industries and face unique computing challenges – whether improving yields for semiconductor wafer inspection, increasing throughput in medical imaging, rendering high-quality animation within a defined budget, or packing enormous processing capacity in deployed ground vehicles. Why are we so driven to tackle these difficult computing problems? Because at Mercury your challenges drive our innovation.

Our new PowerBlock™ 200 system incorporating Cell technology delivers unprecedented processing density within a rugged ATR chassis for vetronics applications. Contact us to learn how our products and services can optimize your challenging applications.

Let us design an innovative solution for you.

www.mc.com/rtc3 | **866-627-6951**

Computer Systems, Inc.
MERCURY
Challenges Drive Innovation™

***RTC*: Product differentiation. First off, how does VersaLogic differentiate itself from its competition being that most makers of PC/104 hardware follow the standard? Second, how do PC/104 vendors differentiate their products from others' small form-factor products in the embedded space such as 3U CompactPCI and CompactPCI Expres,s and forthcoming AMC and MicroTCA?**

Crane: The differentiation between PC/104 vendors probably has little to do with following the PC/104 standard. As with other businesses, the shape, size or compliance of the product is only a small part of what is being sold.

Other factors that matter to a customer include the reliability and efficiency of the circuit design, the component quality, the manufacturing quality system and degree of testing that the product must pass, the service and technical support that's provided, the production scheduling and manufacturing management that assures on-time delivery, the number of emergency spares that are kept in stock, etc. I suppose I should even mention price in there somewhere. In case you haven't guessed, VersaLogic emphasizes quality and services, as do many other PC/104 vendors, rather then per-piece price.

Of course, not all PC/104 product lines are identical. There are 586 to Pentium M processors, plain to fancy analog I/O, parallel and serial I/O cards of all types, AC/DC power supplies, enclosures, extended temperature, low cost, high integration, extremely ruggedized, conformally coated, etc. A big part of what differentiates a company is its experience and expertise in the type of PC/104 products that it offers. Again, the breadth of support for the PC/104 format allows customers to draw from a variety of organizations to get all the different pieces and exactly the expertise that they need for a given application.

PC/104 differentiates itself from some of the current and emerging form-factors primarily by size. The size of the board is small, and the number of PC/104 manufacturers is large. The small size of the PC/104 board seems great, until you realize that there's only so much that you can fit on there, especially with the connector space that's occupied on both sides of the board. For higher functionality or higher performance, you must move to larger platforms. On the other hand, the compact stackable PC/104 format provides a highly expandable cost-effective system that doesn't require a card cage or external interconnect system. What really differentiates PC/104 from other platforms is each specific application and what kind of fit there is. No one platform is appropriate for every application.

Given its long life, high flexibility and wide industry support, the PC/104 and PC/104-Plus platform is going to remain the perfect solution for number of applications. It will be difficult for any other form-factor to displace it in the foreseeable future.

See VERSALOGIC at the RTECC Greenbelt

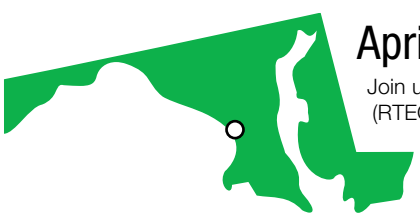

April 25, 2006

Martin's Crosswinds Conf. Center

Join us at the Real-Time & Embedded Computing Conference (RTECC) to learn why quality and services set VersaLogic apart from its competition–and how they differentiate their products from others' small form-factor products such as 3U CompactPCI and CompactPCI Express, AMC and MicroTCA. Find out first-hand about the reliability and efficiency of the circuit design, component quality, manufacturing quality system and degree of product testing, and the service and technical support that's provided. Go to www.rtecc.com/greenbelt to view the open-door technical conference sessions, luncheon keynote, exhibitor list and pre-register to have a badge waiting when you arrive.

Real-Time & Embedded computing conference

www.rtecc.com

PCI Express Connectivity
10Gb per Second
Imagine The Possibilities...

MAX EXPRESS™

One Stop Systems provides complete PCI Express expansion solutions for motherboard, SHB (System Host Board), and CompactPCI Express architectures.

ONE STOP SYSTEMS
2235 Enterprise, Suite 110
Escondido, CA 92029
Tel (760) 745-9883 Fax (760) 745-9824

Call today.
(877) GET-CPCI
www.onestopsystems.com

Evaluating Real-Time Operating Systems for DSP

An RTOS for a DSP processor must meet a different set of system requirements than one designed for a control processor.

by Tom Barrett
Quadros Systems

DSP applications are becoming more complicated with the addition of system control and communications functions formerly relegated to microcontrollers. Yesterday's monolithic assembly language programs cannot meet these demands, particularly with today's engineering schedules.

To face these changes, developers are choosing tools and RTOSs optimized for these evolved DSP architectures. And, this involves a shift. Many DSP-based applications must be designed to run under the control of a real-time operating system (RTOS) to meet demands for system flexibility, code reuse, and rapid project development and deployment.

But, when it comes to DSP, not all RTOSs are created equal. Most traditional RTOSs have their roots in microprocessor or microcontroller designs and offer neither the architecture nor the performance needed for DSP applications. The traditional control-based RTOS architecture can work in a DSP application, but in a less than efficient manner because it is not designed to provide the kinds of code scheduling and processing models required by DSP. Rather than struggle with such compromises, it's often wiser for the developer to choose an RTOS specifically designed, tuned and adapted to the application demands of digital signal processing.

The Loop versus Multitasking

The simplest DSP operating system—and it's a stretch to call it an "operating system," is a loop in which CPU control passes from one point to another in the process of looking at enabling criteria. When a process is ready, the loop allows execution of the associated body of code.

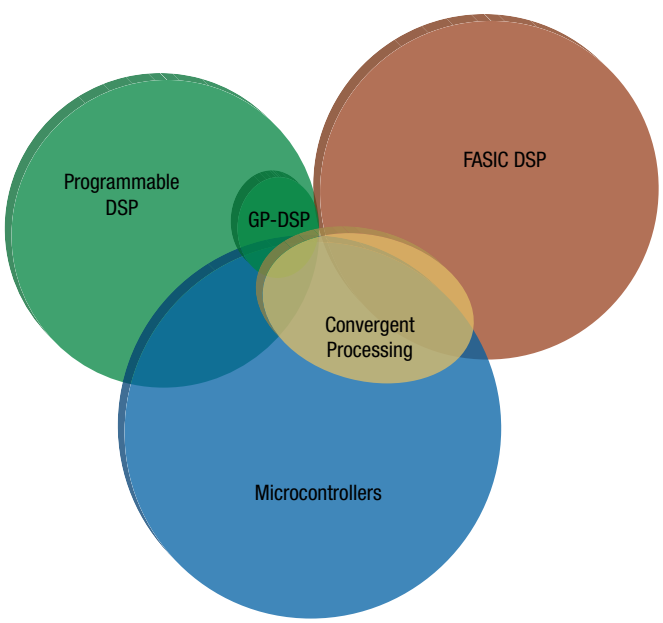

Figure 1 Convergent processing includes elements from traditional microcontrollers, programmable DSPs and signal processing functions built on programmable logic or ASICs.

Get Connected with companies mentioned in this article.
www.rtcmagazine.com/getconnected

In such a design, the operating loop can use a lot of machine cycles, testing the enabling criteria in each pass through the loop. These designs must be handcrafted to preserve application performance. The benefits are minimum code space and speed. The downside is that they are often too well tuned to the particular application, requiring code changes when application requirements change. This results in increased risk and cost.

Contrast this loop structure to the typical microprocessor/controller RTOS that uses a multitasking structure instead of a loop. In a multitasking design, the application is decomposed into several tasks. Each task runs according to a scheduling policy that invokes the RTOS only when it makes a kernel call for some service.

The kernel call may run as a service routine, returning directly to the calling task or it may block the calling task due to the unavailability of some resource or event needed by the calling task. In this case, CPU control passes to another task to run. Whenever there is such a transfer of control, the processor context of one task must be saved before that of the other task can be restored. This process can require many cycles, depending on the processor's programming model. The design allows most of the processor's cycles to be used by the tasks (the application), rather than by the scheduler. The limitation is in system responsiveness and efficiency due to the high overhead of switching between tasks.

So where does that leave the developer? Must he choose between the overhead constraints of a multitasking architecture or the lack of control of a loop architecture? Ideally, a RTOS for a modern DSP processor should be able to apply the advantages of both loop and multitasking behaviors to meet today's DSP application needs.

Convergent Processing

Clearly, with the many negatives and few positives about operating systems for these new DSP processors, it is time to consider a fresher approach to RTOS designs—one that can accommodate the seemingly disparate requirements of digital signal processing and control processing on a single core. It is then appropriate to call the processor that does both types of processing a convergent processor.

An RTOS for a convergent processor incorporates a dual process scheduling model that meets the needs of both DSP and control applications, allowing it to be solely a DSP RTOS, or 100% dedicated to control operations, or anywhere in between the two, hence the term convergent processing (Figure 1).

As an example of how the two models work to meet the demands of their respective domains, look at Table 1, which shows a comparison of operation times required to schedule a DSP process and a control process on a Freescale StarCore processor, an excellent example of a convergent processor.

The results clearly show that the RTOS operations to schedule the DSP process from an interrupt service routine or from another DSP process require significantly fewer cycles, translating into faster response times for those processes versus their control process counterparts.

When it Comes to Performance, Don't Ask the Wrong Question

Unfortunately, DSP developers typically ask the wrong performance question when they evaluate an RTOS. "How much overhead (in MIPS or percentage of CPU loading) does the RTOS take?"

This question stems partly from a mistaken notion of how today's operating systems function. As previously stated, the traditional DSP view of an operating system is a loop that runs continuously looking for a process to perform. When a need arises, it calls the routine associated with the need and executes its code. When the routine is complete, the loop resumes looking for another processing need. Obviously, that design did indeed take extra cycles, and those extra cycles would be a concern to the application developer. Thus the origin of the question about how many MIPS the operating system consumes is understandable.

However, a typical microprocessor/controller RTOS uses a multitasking structure in which tasks get control of the CPU according to some scheduling policy implemented by the RTOS Scheduler. The normal mode, in contrast to the loop method, is for the CPU to be occupied doing application processes. The RTOS kernel call may run as a service routine, returning directly to the calling application process (task). Or, the kernel may block the calling task due to the unavailability of some resource needed by the calling task, in which case the Scheduler will pass CPU control to another task to run. Most of the processor's cycles are used by the tasks. The operating system only takes processor cycles when the tasks call on it to perform some service. Because of this, the RTOS' overhead isn't a fixed value—it depends on how much the application uses RTOS services.

Therefore, a better measure of RTOS overhead is the operational times (or cycles) of the kernel services used by the application.

Measure the Time That Is Important

Another important consideration is the operating system's interrupt response time and the time required to switch from one process to another. Many DSP-based systems depend on interrupts to signal when data should be gathered from or distributed to the outside world and perform the number crunching as fast as possible the rest of the time. In contrast, a control RTOS may have far fewer interrupts. Calculating the response and process switch times lets you determine how many interrupts per second the RTOS can handle at a reasonable processor load.

Of particular importance is the time required to save and/or restore the processor's context. During a context switch in multitasking, the RTOS stores the state of the current task (including the contents of registers) so it can resume the task later. Meanwhile, it restores the registers of the task that is gaining control of the CPU.

The typical DSP has a lot of registers that must be saved and restored during a full context switch. For example, the Freescale MSC8101 DSP processor, based on StarCore IP, has about 68 different registers (about 272 bytes) while Analog Device's Blackfin BF533 has almost 40 different registers. At worst case, saving all of these registers during an interrupt or a context switch can

quickly eat up a measurable percentage of processor cycles if steps aren't taken to mitigate the effects.

On an MSC8101 running at 300 MHz with an interrupt load of 25,000 interrupts per second, simply saving and restoring all 68 registers consumes over 1% percent of the available processor cycles, ((68 registers * 2 cycles per register [save + restore]) * 25,000 Hz) / 300 MHz). This assumes single cycle reads and writes to memory and no processor stalls. Since it's often necessary to do a full context switch while executing control code, the efficiency of the save/restore logic has a great effect on overall system throughput. The optimal configuration of registers to save and restore, and the method to do so, varies with processor architectures and even applications.

Clearly, the DSP application that can't afford time for the save/restore cycle of a multitasking RTOS will benefit greatly from an RTOS design that does not require the save/restore cycle or minimizes its effect.

Architecture Impact

Processor architectures can have other impacts on the operating system besides the number of registers to save. The use of separate memory spaces for I/O and code, segmentation of memory into multiple data spaces and handling of program stacks are all DSP architectural features that the RTOS must accommodate. It isn't uncommon for a DSP processor to have specialized hardware stack support using very fast internal memory. While such stack architectures are efficient for traditional loop-based software designs, they're challenging to the developer of a multitasking RTOS who must efficiently swap stacks for modern preemptive scheduling.

DSPs indeed differ from the typical microprocessor, and the DSP RTOS must adapt and accommodate the architectural quirks of the DSP. These quirks are part of the DSP's architecture and usually exist to make certain operations very efficient. But without the RTOS developer's careful attention to that level of detail, those efficiencies may be lost.

Let it Flow

Another efficiency to consider is how the RTOS handles block data transfers—an essential task in many DSP applications. Typically, a system gathers a block of data and moves it through transform operations (e.g., converting integer data to floating-point data, filtering or compression) before performing the main number crunching.

In handling these transform operations, the RTOS ought to provide facilities that let the application touch the data blocks as little as possible, and with very low overhead. An RTOS that creates unnecessary copies of the data blocks to pass from process to process is a real performance killer. An efficient way of handling blocks is to offer pipes to flow data from one process to another while noting the processor's memory-handling schemes.

With so many concerns about using an RTOS in a DSP-based project, it may seem that it's not worth the effort—a feeling common to experienced DSP developers, charter members of "The Society to Save a MIP." Over the past decade, the microprocessor world has increasingly had to deal with projects too complicated to finish under tight time constraints without looking at the development process. They have shifted from custom to commercial RTOS designs to simplify the process, improve maintainability and evolution, and achieve a high degree of code reuse.

The DSP world today faces the same issues. The essence of the solution is to fit the processor to the application and choose an RTOS that complements the processor capability and the requirements of the application. With proper CPU selection and by scaling the RTOS to meet the application's requirements, the developer will be able to build more efficient and responsive systems, delivering more functionality at a lower cost. ◢

Quadros Systems
Houston, TX.
(866) 879-7892.
[www.quadros.com].

RED ROCK TECHNOLOGIES
Mass Storage Modules
for VMEbus and CompactPCI®
NEW Product!
DVD-RW / CD-RW / CDROM
in VMEbus Form Factor
with Ultra Wide SCSI LVD interface.
See the full line of VMEbus single and
multi-slot mass storage module products at
www.RedRockTech.com
or call Toll-Free: 800-808-7837
Red Rock Technologies, Inc. — 480-483-3777

Comparative Process Scheduling Times	
Schedule DSP process from Control process	225 cycles
Schedule DSP process from an ISR	63 cycles
Schedule DSP process from another DSP process	82 cycles
Schedule a control process from another control process	354 cycles

Table 1 Operation times for control and DSP process scheduling on a Freescale StarCore processor.

COM Express Module Uses Intel Core Duo Processor

A new COM Express CPU module integrates an Intel Core Duo processor L2400, together with the Intel 945GM Express chipset and the matching southbridge ICH7-M. Intel Core Duo dual kernel processors combine the advantages of two high-performance processor kernels with intelligent power management. The Toucan from Lippert Embedded Computers has all interfaces defined by the COM Express specification. These are routed to a custom specific baseboard using suitable connectors. The chipset features, among other functions, fast graphics and display controllers, which drive display devices over either CRT or LVDS. If necessary, an external graphics controller can be connected using the PCI Express (PCIe) x16 lane. More high-speed peripherals may be connected via the four available PCIe x1 lanes. Gigabit Ethernet, serial ATA and IDE are also on board, as well as eight USB 2.0 ports and sound output. Up to 2 Gbyte RAM provide enough memory, even for large applications.

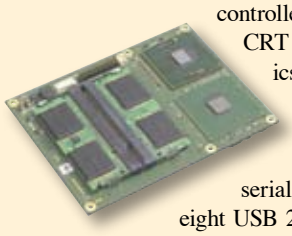

The module has built-in hardware monitoring and status LEDs show the current state of the module. The integrated Intel Active Management Technology (AMT) functions allow remote diagnostics and firmware updates over the network. The Toucan is especially developed for use in extreme environmental conditions (-40° to +85°C, shock and vibration) and does not need to be actively cooled. The Toucan runs standard operating systems like Windows XP and Linux. The RoHS-compliant module is optionally available for the extended temperature range of -40° to +85°C. Quantity 100 pricing starts at $799.

Lippert Embedded Computers, Atlanta, GA. (404) 459-2870. [www.lippert-at.com].

VITA 31 VME Board Targets Rugged Mission-Critical Apps

The CPC600 from Fastwel is a COTS product built for highly reliable mission-critical VME systems. The board is based on an Intel Pentium M processor running at up to 2.0 GHz and supports up to 2 Gbytes of DDR SDRAM with ECC. It is equipped with an Intel 82855GME GMCH & Hance Rapid ICH chipset.

All components including CPU, SDRAM and a 32 Mbyte solid-state disk may be soldered on board thus providing superior shock (50g) and vibration (5g) resistance. Four Gigabit Ethernet ports and conformance to the VITA 31 specification makes CPC600 a platform suitable for robust redundant systems. Live Insertion support, smart temperature control, HW monitor and watchdog timer make CPC600 suitable for mission-critical applications. Software support for this SBC includes MS DOS, Windows 2000/XP/CE, QNX and Linux.

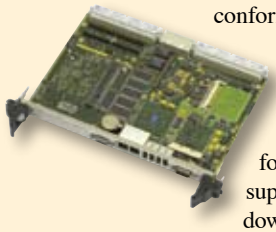

A key element of the crash safety subsystem is 32K non-volatile RAM where the user application can keep critical data and system log, which should be kept even if power fails. CPC600 also has 64K EEPROM memory for user applications. The board has a multi-drive direct mounting feature—it supports 2.5" and 1.8" HDD simultaneously. The board is rated to withstand operating temperatures from -40° to +85°C. Pricing starts at $2,987.

Fastwel, Moscow, Russia +7 (495) 234 0639, [www.fastwel.com].

Microstepping Drive Has Programmable Resolution

Single-axis motion controllers have a basic cost advantage over off-the-shelf, card-based approaches, as well as fewer cables and greater reliability. One such drive is the microstepping version of Performance Motion Devices' ION Digital Drive, which also features microstepping resolution that is programmable up to 256 microsteps/step.

ION is a fully enclosed module that delivers high-performance motion control, network connectivity and amplification for DC brush, brushless DC or step motors. It is targeted toward medical, scientific, semiconductor, automation, industrial and robotic applications. The ION Microstepping version supports two-phase bipolar motors and can automatically switch between user-defined current modes. The drive includes asynchronous serial (RS-485) or CANbus network communications, as well as precise current control.

Selectable profiling modes include S-curve, trapezoidal, velocity contouring and electronic gearing. The ION Microstepping version accepts commands over the communications network, such as position, velocity, acceleration and jerk, and generates a corresponding trajectory on-the-fly. Additional features include 500W capability, 8 Amps continuous current and 15 Amps peak current, a 102 microsecond servo loop rate and field-oriented control. Prices start at $223 in OEM quantities.

Performance Motion Devices, Lincoln, MA. (781) 674-9860. [www.pmdcorp.com].

Virtual Platform Available for Multi-Threaded MIPS32 34K Cores

To design high-performance embedded systems with the MIPS32 family of 34K cores requires an understanding of multi-threading techniques. A new version of the Virtio Virtual Platform, the VPMM-SC, lets development teams perform in-depth evaluations of this technology's impact on their designs.

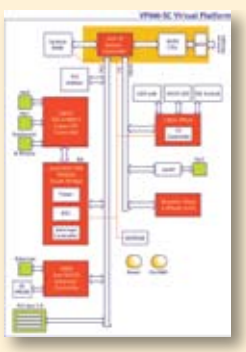

The VPMM-SC Virtual Platform provides fast system simulation of the MIPS32 34K core-based MALTA development board. Developers can quickly reconfigure the MIPS32 34K cores in a variety of different multi-threading configurations and evaluate the system impact with detailed internal core statistics. The platform uses a combination of C-code, graphical modeling and transaction-based modeling to replicate the full functionality of the MALTA board in a PC-based simulation.

The VPMM-SC supports the MIPS32 4KE, 24K and 34K core family members and offers a choice of instruction-accurate and cycle-accurate CPU models. The VPMM-SC supports standard operating systems such as Linux. It can execute binary code files that target the MALTA board and integrates with essential software development tools such as Green Hills MULTI and GDB debuggers. Pricing starts at $1,990.

Virtio, Campbell, CA. (408) 341-0844. [www.virtio.com].

VMEbus SBC Provides Dual PowerPC 7447s at 1GHz

Targeted at high-end applications for defense, aerospace and compute-intensive operations, a new VMEbus-based single board computer (SBC) also provides a set of PC system functions that make it a self-contained single board computer with workstation features. The Maverick from General Micro Systems has two 1 GHz 7447 PowerPCs with 128-bit AltiVec technology, on-die 512 Kbytes L2 caches, up to 1 Gbyte of DDR SDRAM provided by SODIMM modules, full Symmetric Multi-Processing support and the Discovery-III system controller. Additionally, 32 Mbytes of application flash and 2 Mbytes of bootable flash let users store reasonable-size programs on board, eliminating the need to download.

Two 64-bit, 66/100 MHz PMC expansion sites—one with rear I/O—offer design flexibility for customization to specific needs. This makes the board suitable for uses such as image processing, signal processing, data analysis and high-end machine control, as well as for applications in radar and sonar systems and systems used for training, simulation and targeting. The Maverick incorporates dual USB, IDE DMA-100, Serial, PS2 mouse and keyboard and ultra SCSI 320/160 with differential output. For flexibility in setup and operation, the V394 offers three Gigabit Ethernet ports, one Multi-Protocol Sync/Async Serial port, parallel port, 21 Discreet Digital I/O lines, and front panel LEDs for POST code and Status display. Support for Linux, QNX and VxWorks operating systems is standard. Pricing starts at $2,600 for quantities of 100 units.

General Micro Systems, Rancho Cucamonga, CA. (909) 980-4863. [www.gms4sbc.com].

AMC Module Runs Pentium M at 2.0 GHz

A processor module suitable for AdvancedTCA, MicroTCA and proprietary platforms, the PR AMC/33x from Concurrent Technologies is a single-width, full-height AMC processor module supporting the 2.0 GHz Intel Pentium M processor 760. The PR AMC/33x also supports the Intel E7520 server chipset, which interfaces up to 2 Gbytes soldered DDR2-400 ECC memory giving a peak memory bandwidth of 3.2 Gbytes/s. To provide a higher MTBF and improve cooling, all components are soldered including the memory and processor. The PR AMC/33x is designed in compliance to AMC.0 (including full hot swap and IPMI capabilities), AMC.1 Type 8 (PCI Express x8), AMC.2 Type E2 (2x Gigabit Ethernet) and AMC.3 Type S2 (2x Serial ATA150 ports). The module also features two USB 2.0 ports, two RS-232 ports and 1 Mbyte application flash.

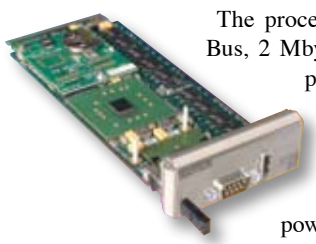

The processor supports a 533 MHz Front Side Bus, 2 Mbytes L2 cache with an enhanced data pre-fetcher, merges CPU operations with micro-ops fusion, and has advanced instruction prediction of program branches, dedicated stack manager in hardware, intelligent power distribution (clock gating based on program demand), Deep(er) Sleep Alert States and Enhanced Intel SpeedStep technology. The Intel E7520 includes three 8x PCI Express links and a DDR2-400 memory interface to help increase I/O bandwidth and reduce system latency for data-intensive applications. Supported operating systems include Linux, Windows 2000/ Server 2003/XP/ XP Embedded, QNX and VxWorks. Pricing starts at $1,735.

Concurrent Technologies, Colchester, UK. +44 1206 752626. [www.cct.co.uk].

M2M Communications Terminal

By now, a lot of industrial equipment has been connected to the Internet. But not all M2M communications terminals are equal in capabilities, features and cost. The eDindus terminal from eDevice handles most of the tasks involved when connecting industrial equipment to the Internet over GPRS, PSTN (V90) and Ethernet networks.

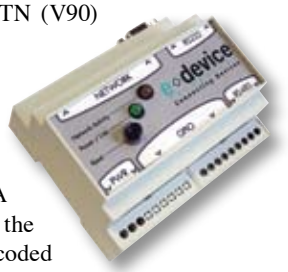

The eDindus includes an FTP client, an e-mail client and a customizable Web server. The terminal's software development kit (SDK) allows scripts to be created and embedded in an interpreted language. A 16-bit ATMega microprocessor manages the Modbus protocol and allows specific C-coded client modules to be embedded in the terminal. The internal eDmod module manages network connectivity, while an Atmel 16-bit microcontroller runs I/O and C-code components.

eDindus can be connected to any equipment through a wide range of hardware interfaces, including RS-232/ RS-422/RS-485, analog and digital input/output and dry switch. The terminal is packaged in a rugged DIN-rail plastic casing with screw connectors. Additional connectors include nine external general-purpose I/O software-configurable as digital I/O or a 10-bit A/D converter and 50 internal general-purpose I/O with HE14 connectors. The terminal features 2kV galvanic isolation and an 8 to 60VDC power input, with an optional 110 to 220V power supply. Pricing is $240 for the eDindus300 with GPRS Internet module.

eDevice, Merignac, France. +33 (0) 5 56 12 77 79. [www.edevice.com].

Modules Offer High-Speed Simultaneous Measurement via USB

Simultaneous analog input multifunction data acquisition modules for USB 2.0 are now available in the form of Data Translation's DT9832 Series. The DT9832 Series offers two or four channels of simultaneous, 16-bit analog measurement at throughput speeds up to 2 MHz per channel. This series offers very fast simultaneous sampling of four independent analog channels addresses the needs of customers who are attempting to correlate physical phenomena at the same exact instant in time. Now they can do this with fast frequency component signals as found in semiconductor device characterization, nanotechnology testing, scientific analysis, drug discovery and automotive testing.

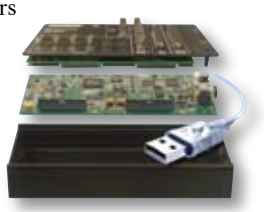

Each analog channel in the DT9832 Series has its own 16-bit A/D converter that utilizes a common clock and trigger for simultaneous and synchronous sampling of all analog and digital inputs at the exact same instant in time. Throughput per channel for the DT9832 is 1.25 MHz on four 16-bit analog inputs. Throughput for the DT9832A is 2.0 MHz on two 16-bit analog inputs. Two 32-bit counter time channels offer flexible clocking and triggering and three quadrature decoders allow determination of X/Y positioning and rotation. All subsystems operate synchronously at 48 MHz. Pricing for both versions starts at $1,970.

Data Translation, Marlboro, MA. (508) 481-3700. [www.datatranslation.com].

Get Connected with companies and products featured in this section.
www.rtcmagazine.com/getconnected

CompactPCI Core Duo CPU Module Also Offers CompactPCI Express

The F15 single board computer from MEN Micro can occupy a single 32-bit/33 MHz system slot on the CompactPCI bus or be used as a stand-alone single Eurocard. The module is powered by the new Intel Core Duo dual-core processor and is available with either the 1.6 GHz, 15W version or the 2.0 GHz 31W version. The board also uses the new Intel 945GM chipset, which allows for six PCI Express and two SATA lines for communication needs besides the conventional PC functions. The F15 can be used as a system slot board in CompactPCI using the J1 andJ2 bus connectors. Alternatively, it can be fitted with a CompactPCI Express-compatible side card in a hybrid backplane that supports CompactPCI Express. In the latter case, the side card's newly defined ADF-HF and eHM connectors route some of the available PCI Express lanes to the backplane.

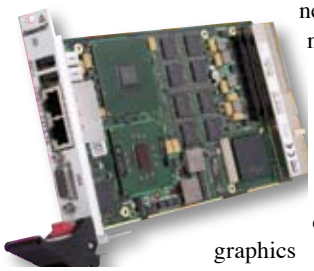

Memory consists of up to 2 Gbytes of DDR2 DRAM plus a CompactFlash slot that can alternatively be equipped with a 1.8-inch hard disk. Front panel I/O includes a VGA graphics connector, two PCI Express-driven Gigabit Ethernet interfaces and two USB 2.0 ports. The F15 can also be fitted with different side cards for use in 6U systems and the incorporation of additional I/O functions. The board comes with board support packages for Windows, Linux, VxWorks and QNX operating systems. Single unit pricing starts at $2,154.

MEN Microelectronics, Lago Vista, TX. (512) 267-8883. [www.menmicro.com].

COM Extends Design Life of ETX-Based Systems

A new XTX 820 Computer-on-Module (COM) from Ampro Computers gives OEMs access to the latest chipset technology without the need to completely redesign their ETX baseboards. Based on a Pentium M with performance of up to 1.8 GHz, the XTX 820 COM retains legacy peripherals while offering the latest high-speed interconnect and storage technologies.

The XTX 820 COM plugs into a baseboard to allow customization of board size and shape, I/O circuitry, connector locations and easy migration to higher performance modules. It includes 1 GHz, 1.4 GHz or 1.8 GHz processors, the Intel 915GM chipset and up to 1 Gbyte of DDR 400 RAM. For expansion and connectivity, the COM features six USB 2.0 ports, EIDE and SATA interfaces, 10/100 Ethernet, PCI expansion and RoHS compliance.

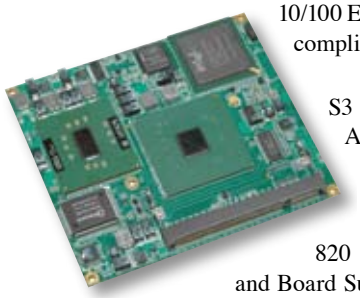

Full ACPI 2.0 support, including S3 suspend-to-RAM, is included. AMI BIOS provides full support for wake-up devices including LAN, keyboard, mouse, power button and PCI or SMBus activity. XTX 820 QuickStart Kits include drivers and Board Support Packages (BSPs) for Windows XP, Windows XP Embedded, Windows CE 5.0, VxWorks and QNX, and a full Linux 2.6 distribution (Fedora Core 3). Prices start in the low $500s in moderate quantities.

Ampro Computers, San Jose, CA. (408) 360-0200. [www.ampro.com].

Backplane Implements New VXS Processor Mesh Architecture

The VXS Processor Mesh architecture recently proposed to VITA delivers 112.5 Gbytes/s of aggregate throughput within the processing mesh in a single chassis, a 6x improvement over currently available technology. The industry's first 12-slot backplane that implements this architecture is now available from Elma Bustronic.

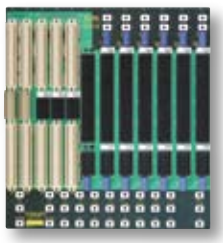

This hybrid backplane was developed to enable a switch/processor mesh technology for applications requiring multiple boards for application processing. It implements two VME64x slots, three VME64x/VXS payload slots and six VXS switch slots. Each switch slot implements twenty x4 links for a total of 25 Gbytes/s per switch slot. The system architecture supports throughput of up to 7.5 Gbytes/s between the I/O front-end and the processing mesh.

Other features include a rear-facing I/O slot that conglomerates all of the payload slot rP0 reserved pins, 16 pins dedicated to user I/O from each of the processor slots and two channels served by the dual-dual star VXS fabric switch. Pricing begins at $3,000, depending on configuration.

Elma Bustronic, Fremont, CA. (510) 490-7388. [www.elmabustronic.com].

Four Embedded Platforms Based on New Intel Core Duo

Kontron is introducing four embedded computing platforms with new Intel Core Duo dual-core processors: a mini-ITX motherboard, a 3U and a 6U CompactPCI board, and an ETXexpress module, which meets the COM Express specification in accordance with PICMG COM.0. Each Kontron board is equipped with the 2 GHz Intel Core Duo processor T2500; three of them with the Mobile Intel 945GM Express chipset and the 6U CompactPCI board with the Intel E7520 and 6300ESB. They all offer up to twice the performance of the Pentium M processor at the same GHz clock speed. Power consumption, at up to 31W, is extremely low. These new boards are expected to make integration of two or more previously stand-alone applications into a system possible without any interference between them. These new processor features can help significantly reduce hardware costs and time-to-market for OEMs. Applications such as visualization and control will also be able to be implemented in one system with various operating systems, while maintaining hardware costs under control. Kontron will introduce a total of more than ten different embedded computing platforms with these new Intel Core Duo processors to the market in 2006.

The Kontron mini-ITX board with dual core is targeted at classic IPCs as well as interactive gaming machines with two monitors for two players. Its essential advantages over commercial boards are its long-term availability and low power consumption. The Kontron 3U CompactPCI with dual-core processor is aimed at performance-hungry, high-availability and robust industrial applications. The 6U CompactPCI board with IPMI brings higher CPU density primarily in the telecom and datacom sector as well as in the high-end medical and image processing sector, helping in saving space. The ETXexpress modules offer dual-core performance for customer-specific embedded designs. Pricing starts at approximately $256 for the mini-ITX (plus processor), $1,000 for the ETXexpress module and between $2,400 and $3,150 in OEM quantities.

Kontron America, Poway, CA. (888) 294-4558. [www.kontron.com].

Get Connected with companies and products featured in this section.
www.rtcmagazine.com/getconnected

Parallel DSP Board Offering 72 GIPS of Processing Power

A new parallel DSP board provides over 2000 Mbytes/s of concurrent onboard I/O and the ability to combine any number of customer-specific channels of data for up to 500 Mbytes/s of continuous data input from an external source. The PMP1000 from Signatec harnesses up to nine Texas Instruments TMS320C6414T DSPs. The board maps 64 Mbytes of memory directly into each of the eight processing DSPs, and incorporates a parallel DSP operating system. Operating at a clock rate of 1 GHz, each DSP is capable of executing up to 8 instructions per clock cycle for a peak processing performance of 8 GIPS per DSP, or 72 GIPS for the entire board. A four-DSP option also available.

PMP1000 includes three data interfaces. The primary interface is the PCI-X host interface, which can transfer data at peak rates of up to 1 Gbyte/s when the board is installed in a PCI-X slot with 64-bit data at 133 MHz. The board can also be operated in lower performance PCI slots, including both 64 and 32 bits from 133 MHz to 33 MHz. The second is the Signatec Auxiliary Bus (SAB), which can sustain transfer rates up to 500 Mbytes/s with other Signatec products. This interface supports up to eight boards. The third external bus interface allows for high-speed external interface connections up to 500 Mbytes/s to any customer-specific bus or combination of buses. Pricing starts at $8,400 for the four-DSP version.

Signatec, Corona, CA. (951) 734-3001. [www.signatec.com].

Hybrid FPGA/PowerPC Processor Is Fully Rugged

In applications demanding increasingly powerful front-end signal processing, lots of compute power and ruggedization are needed. A new hybrid FPGA/PowerPC processor card from Radstone is available in four ruggedization levels.

The V4DSP features dual Virtex-4 FX 60/100 FPGA processing nodes combined with a Freescale MPC7448 PowerPC processor operating at speeds of up to 1.4 GHz.

The MPC7448 PowerPC's AltiVec general-purpose processing node can be used as the first back-end application processing stage to reduce chassis slot count. The card has two StarFabric ports at the P0 connector to enable multiboard connectivity. For mezzanine support, either a 64-bit/66 MHz PMC site can be configured, or an XMC site equipped with eight multigigabit transceivers to the two Virtex-4 nodes, each operating at up to 3.125 GHz.

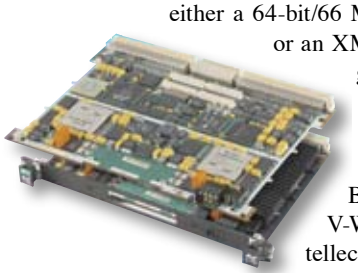

Included are a Radstone Board Support Package (BSP) and V-Wrap, a comprehensive FPGA intellectual property (IP) wrapper, which is Radstone's pre-prepared set of VHDL interfaces to the Virtex-4's embedded I/O and memory. Also included is support for Radstone's acclaimed AXIS Advanced Multiprocessor Integrated Software environment for developing highly scalable multiprocessor systems. The V4DSP is available in four ruggedization levels. Price for a single, level-one unit is $13,606 in OEM quantities.

Radstone Embedded Computing, Towcester, UK.
+44 (0) 1327 359444. [www.radstone.com].

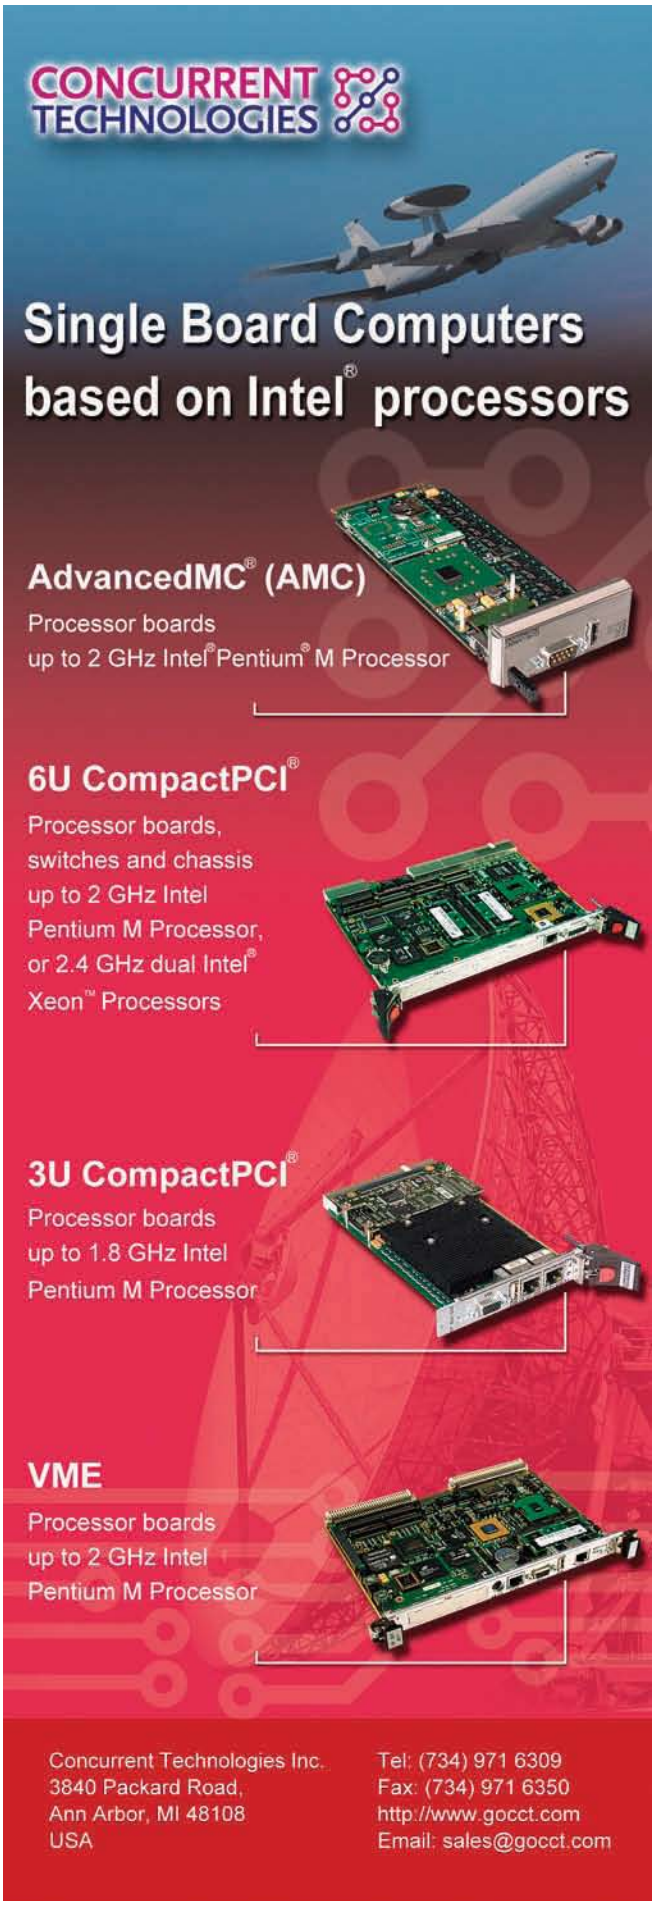

MISSION CRITICAL VME/cPCI
data storage modules

Extreme Comprehensiveness: We offer the most comprehensive VME/cPCI storage product line in the world, offering device alternatives for any standard or unique application.

• Solid State Disk • Removable Hard Disk
• Tape Drives • Optical Disk • PCMCIA Adapter

Extreme Performance: Our VME products feature extreme speed, capacity and ruggedly reliability with 320 MB/sec throughput enabled by LVD SCSI technology, storage capacity of more than 600 GBs per module and a 1,400,000 hour MTBF.

Extreme Quality: Phoenix International is the only manufacturer of VME data storage products that is ISO 9001:2000 Certified.

PHOENIX INTERNATIONAL

Phoenix International Systems, Inc. An ISO 9001:2000 Certified SDVOSB
714-283-4800 • 800-203-4800 • www.phenxint.com

Supporting Your Future HUNT-ENGINEERING

USB connected Programmable FPGA systems

Data Acquisition
- 2 ch 210Msps 12 bit A/D
- 2 ch 160Msps 16 bit D/A
- Configurable digital I/Os
- Virtex-II FPGA (1.5M or 3M)
- USB 2 or Standalone

Configurable Digital I/O
- Virtex-II FPGA 1M gates
- 90 Configurable Digital I/Os
- RS232/485 and LVPECL
- USB 2 or Standalone

Programmable hardware with cables, device drivers, loading tools, examples and Power Supply.
Systems can be used connected to a PC using USB, or can function standalone (without USB) using the initialisation PROMs.

sales@hunteng.co.uk, +44 (0)1278 760188
www.hunt-rtg.com

Carrier Card Delivers 10X Boost in Image Processing Speed

When designing complex, high-performance image processing subsystems for advanced industrial inspection, designers want to increase performance and reduce system size. With that in mind, TEK Microsystems has introduced a carrier card that can reduce board count by up to 40%, as well as slashing weight, power and volume.

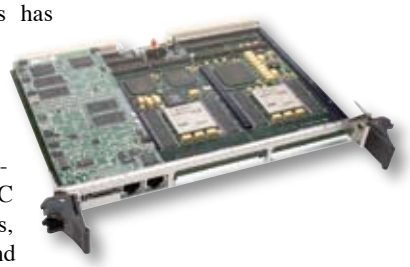

The PowerRACE-3A combines high-density PowerPC and FPGA processing engines, an onboard switched fabric, and PMC interfaces. High speed I/O and image processing are done in a single slot, while FFT, pulse compression and image processing are handled on PMC carriers. The card uses two 800 MHz 440GX PowerPC processors to support high throughput without incurring host processor overhead. An onboard fabric allows each PMC site to transfer data concurrently to off-board RACE++ ports, FPGA processing or memory, eliminating fabric contention and maximizing overall system performance.

Included with the PowerRACE-3A is the tekX software environment, which provides tools for fabric configuration, buffer management, data transfer, interprocessor communications, data storage/playback and integration of streaming FPGA and I/O modules. The PowerRACE-3A is available 10 to 12 weeks ARO. Pricing begins at $17,995.

TEK Microsystems, Chelmsford, MA. (978) 244-9200. [www.tekmicro.com].

OPTO Expansion Card Delivers I/O Voltage Isolation

High-performance, I/O-intensive industrial control applications often require high-voltage isolation. The flexible I/O base address range and 16-bit pass-through connector in a new user-configurable EPIC-sized OPTO expansion card from Micro/sys is designed for rugged, industrial systems with analog and digital I/O that require high-voltage isolation.

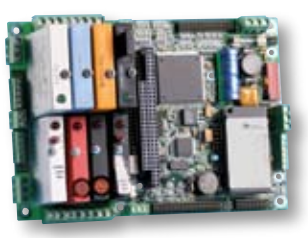

The OPTO104 accommodates eight industry-standard digital or analog OPTO plug-in modules, as well as PC/104 CPU or I/O cards. Any combination of digital or analog I/O modules can be plugged into the 4.53-in. x 6.50-in. board. For applications demanding more I/O, up to four OPTO104 boards can be connected in a slave configuration using a 26-wire ribbon cable. This produces 32 available slots for OPTO 22 G4 modules for AC or DC digital I/O, and/or Grayhill G5 modules for analog I/O. All of these modules offer opto-isolation up to 4000V.

The board operates at 5V ± 5%, with an extended operating temperature range of -40° to + 85°C. Alternately, DC/DC converter options allow input voltages from 9V to 75V. A user-selectable jumper can bleed off common mode leakage current to the negative inputs or to safety ground. Price is $375 in single quantity.

Micro/sys, Montrose, CA. (818) 244-2600. [www.embeddedsys.com].

Get Connected with companies and products featured in this section.
www.rtcmagazine.com/getconnected

The Power of Convergence...

Reliability. Scalability. Mobility...on a gateway blade.

Open architecture drives costs down and performance up.

Modular VoIP "blades" based around standard, interoperable modules like PMC and AdvancedMC reduce costs by limiting the number of unique blades that telecom OEMs and carriers have to purchase and stock. A softswitch or media gateway controller can be deployed in a minimal configuration and scaled up later (to OC-3 and beyond) without replacing the whole blade and without taking it offline. SBE provides high-performance DSP resource modules that deliver premium carrier class voice processing with world-class features using Texas Instruments' DSPs with Telogy Software. In addition, these modules support transcoding and transrating to enable the integration of voice, video, data, and wireless.

SBE products are scalable from daughterboard modules to complex gateway blades, and provide telecom carriers/service providers with a choice of programmable voice platforms featuring SBE's line of network interface cards, ranging from T1 and T3 to Gigabit Ethernet and IPsec/SSL/WLAN acceleration. Full Linux support is available on every board.

sbe

Linux On Demand

flexibility on demand | 925-355-2000 | info@sbei.com | www.sbei.com

RTC delivered straight to your inbox

Switch to Digital Delivery

RTC magazine now comes straight to your Inbox! Just visit RTCmagazine.com to convert your print magazine subscription to Digital Delivery. Each month your downloadable issue will include everything in the print magazine in a searchable PDF format. Archive your issues for portability and easy reference!

Check out all of *RTC*'s Digital Services

RTC Online Edition

Visit RTCmagazine.com and browse the latest issue of *RTC* magazine online. You'll find all the keen editorial insight, detailed technical articles and product info that make the print edition a must-read for industry insiders.

RTC Update Email Newsletter

Get the latest headlines and product news from RTC Update. The newsletter arrives in your Inbox mid-month, keeping you up to date and on top of our fast moving industry. Sign up at RTCmagazine.com today!

RTC Get Connected!

RTC's Get Connected will hook you up with the right people without a lot of hassle. When you need important product details right from the source, look for the Get Connected symbol in *RTC*'s print and digital versions.

Get Connected

RTC Supplemental Content

Looking for more info? *RTC* Supplemental Content has it. Select *RTC* magazine articles include additional content posted online. Dig a little deeper into the technologies shaping the future of our industry.

RTCmagazine.com

RTC

the RTC Group

PMCs: They're the Real Thing

Since its inception, the PMC has added power and flexibility, first as easily configured I/O, up to and including switched fabrics, and later to adding processor power. The PMC continues to evolve today, giving developers ever-increasing options.

by David Compston
Radstone Embedded Computing

The PCI Mezzanine Card is rapidly achieving prominence as the key element in any system solution. It delivers high performance and wide-ranging functionality in a form that is cost-effective, enables easy upgradeability, is small and lightweight, consumes little power and enables system designers to employ cutting-edge technology in systems that may have been conceived several years ago. Importantly, the PCI Mezzanine Card (PMC) is an industry standard, its specifications defined by IEEE, and as such conforms with then Defense Secretary William Perry's mandate that military buyers—an influential element of the total market—should choose commercial off-the-shelf (COTS) solutions where possible.

It is a standard that is widely supported, with an enormous range of products available—meaning that competition is intense, and prices are low. It is capable of being ruggedized such that it can be deployed, if necessary, in harsh environments. And, best of all, the whole philosophy of the PMC revolves around modularity, upgradeability, technology insertion and the ability to seamlessly boost the performance and capability of legacy systems.

Although the standards for PMCs were not agreed upon until the middle of 2001, product had been shipping to the marketplace for over five years. As a concept then, the PMC has been around for a decade or more—and it's no surprise to hear industry commentators opine that, over that period, several thousand different PMC devices have been created. What

Industry Standards Relevant to PMCs

Reference	Defines
IEEE 1386 CMC	Form-factor, connector, electrical interconnects
IEEE 1386.1PMC	Mapping of PCI bus signals
VITA 20-2001	Conduction-cooled PMC
VITA 32	Processor PMC
VITA 39	PCI-X
VITA 42 (XMC)	High-speed I/O and fabric support
VITA 46	New 3U and 6U form-factors for advanced applications

Table 1 VITA has played a key role in defining standards for PMCs.

Get Connected with companies mentioned in this article.
www.rtcmagazine.com/getconnected

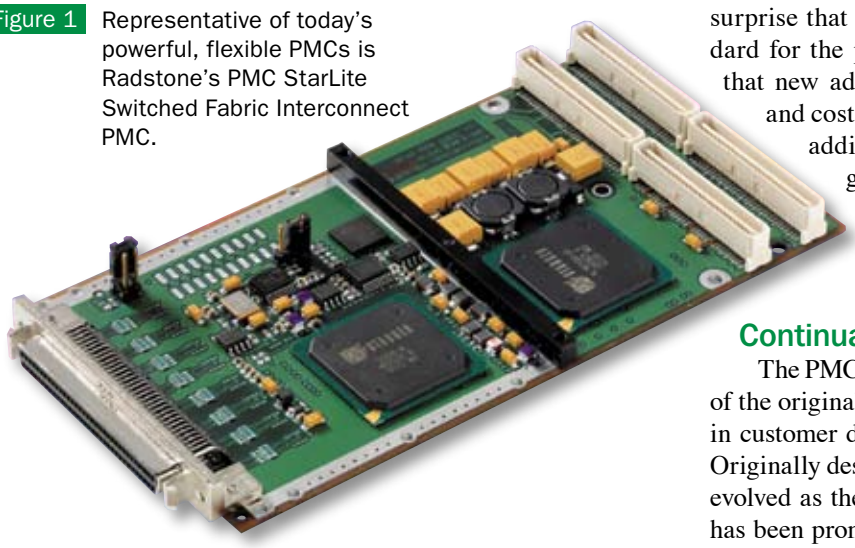

Figure 1 Representative of today's powerful, flexible PMCs is Radstone's PMC StarLite Switched Fabric Interconnect PMC.

has been interesting to follow has been the development of the PMC and the way in which it has been capable of responding to changing customer needs—and the extent to which, today, there is a case for saying that the majority of the key functionality in any system is delivered, not by the carrier card or host, but by the mezzanine card.

In fact, it is probably true to say that the PMC was originally created by both the VMEbus and CompactPCI communities in order to provide more flexible I/O capabilities than was otherwise possible. While applications for PMCs have broadened since that time—the advent of the processor PMC, for example, together with graphics mezzanines and mezzanines for applications such as software defined radio—their primary purpose continues to be in the field of providing off-board communications, with today's products giving existing systems access to the brave new world of serial switched fabrics.

The Role of VITA

The VMEbus community has continued to play a leading role in the development of the PMC, notably via VITA—the VMEbus International Trade Association (Table 1). VITA's first contribution was the VITA 20-2001 standard, which defined the conduction-cooled PMC (CCPMC) that is an absolute requirement for harsh environment applications. Providing the necessary thermal interfaces to a similarly conduction-cooled carrier board, the CCPMC also features extended provision for environments in which shock and vibration are present.

A natural outflow of the increasing focus on COTS solutions—among military customers, for example—was that it would, in theory at least, allow access to successive generations of increasingly powerful microprocessors that would appear with astonishing rapidity, promising the performance enhancements that were necessary to keep the capability of systems in line with the demands being placed on them. The wholesale replacement of entire systems on a regular basis is, of course, unthinkable for the majority of customers—not only for cost reasons, but also for reliability/stability reasons—and so it was no surprise that VITA was also instrumental in creating the standard for the processor PMC (PrPMC). Here, the concept was that new advances in processor technology could be simply and cost effectively leveraged by adding power through the addition of a mezzanine. In fact, the VITA 32 standard goes further in that it allows the mezzanine-based processor to be either the "monarch" (the main PCI bus processor) or to act as a slave processor in a loosely coupled multiprocessing system.

Continual Evolution of the PMC

The PMC also continues to demonstrate the inherent strength of the original concept in its ability to adapt to changes, not only in customer demands, but also in the computing infrastructure. Originally designed around Intel's PCI architecture, the PMC has evolved as the PCI architecture has evolved—and, again, VITA has been prominent in defining (through the VITA 39 standard) how it has done so. PCI-X 2.0 is capable of eight times the bandwidth of PCI and thus of enabling substantial increases in performance, and VITA 39 defines PCI-X operations for both PMCs and PrPMCs.

Originally devised to enhance the range of I/O options available for VMEbus systems, the most recent developments in the PMC world have seen it returning to its roots, with the VITA 42 standard (also known as XMC), which defines how a PMC will support very high-speed I/O and, notably, the serial switched fabrics—such as PCI Express, RapidIO and InfiniBand—that are attracting significant attention as a route toward very high-performance, immensely scalable solutions. The opportunity presented by serial switched fabrics is also addressed by the VITA 46 standard, which is rapidly nearing approval, and which will provide even greater opportunity for VMEbus systems designers to leverage their benefits (Figure 1).

The increasing focus of many applications on communications is serving to heighten the key role of the PMC: the network is, to all intents and purposes, the computer. In communications-centric applications, intelligence and processing power are widely dispersed throughout the total system. There is also increasing focus on compact and lightweight hardware solutions—attributes that the PMC has in abundance. A classic example of the type of application for which PMCs are destined is in the military, where unmanned vehicles are characterized by a development path that sees them becoming smaller, lighter and thus much more widely deployed. Moreover, in these physically constrained environments, power dissipation is a key issue. Here too, the PMC, in its PrPMC form, represents an attractive solution because of its performance/power ratio.

The PMC Today

If "flexibility" has been the watchword in terms of the development of the PMC, the most recent developments are serving only to redefine what that means. Today's mezzanine cards present the systems designer with an enormous range of options. For manufacturers such as Radstone, commercial success lies as

much with the range, power and functionality of the PMCs the company can offer as it does for "higher profile" products such as single board computers, graphics processors and digital signal processors. PMCs give customers access to 1553 communications, high-speed serial and parallel interfaces, software defined radio, Ethernet, flash memory, serial switched fabrics, graphics and CAN Bus functionality—functionality that transforms a box of electronics into a solution.

It is probably an overstatement to say that VITA 32 and the concept of a processor PMC defined a future in which the carrier board could be considered little more than a host. However, in today's architectures, availability of a high level of processing power is a given from the carrier card: the rapidly emerging key functionalities for today's applications are delivered by the PMC. As such, there is a strong element of truth in the statement.

Systems design has moved away from the monolithic architectures of the past, which were relatively inflexible and highly resistant to upgrade, and toward architectures, which are modular and easily—and cost-effectively—capable of being enhanced in line with the changing requirements of a dynamic marketplace. The PMC, in effect, enables a "future proof" system to be designed from the outset, without compromising initial performance. Even better: it allows the unforeseen to be comprehended. Who, for example, in 1995 could have predicted serial switched fabrics? Beyond this, the standards that govern the PMC ensure maximum interoperability, providing systems designers with the opportunity to select at will from the products of a broad range of manufacturers, mixing and matching hosts and mezzanines to deliver the optimum—and most cost-effective—solution. That all of this is possible is largely due to the underlying strength of the concept and the increasing functionality and sophistication of the humble PCI Mezzanine Card. ◢

Radstone Technology
Billerica, MA.
(978) 671-9490.
[www.radstone.com].

Event Calendar

03/28/06
Real-Time & Embedded Computing Conference
Albuquerque, NM
www.rtecc.com/albuquerque

03/30/06
Real-Time & Embedded Computing Conference
Phoenix, AZ
www.rtecc.com/phoenix

04/03-06/06
Embedded Systems Conf.
San Jose, CA
www.commdesignconference.com

04/03-06/06
Linux World Conf & Expo
Boston, MA
www.linuxworldexpo.com

04/03-06/06
22nd Nat'l Space Symposium
Colorado Springs, CO
www.spacesymposium.com

04/03-07/06
CTIA Wireless 2006
Las Vegas, NV
www.ctiawireless.com

04/17-20/06
Annual Nat'l Logistics Conf & Exhibition
Miami, FL
www.ndia.org

04/19-20/06
AFCEA Intelligence Symp.
Washington, DC
www.afcea.org

04/25/06
Real-Time & Embedded Computing Conference
Greenbelt, MD
www.rtecc.com/greenbelt

04/27/06
Real-Time & Embedded Computing Conference
Boston, MA
www.rtecc.com/boston

If your company produces any type of industry event, you can get your event listed by contacting sallyb@rtcgroup.com.
This is a FREE industry-wide listing.

Advertiser Index

Get Connected with technology and companies providing solutions now

Get Connected is a new resource for further exploration into products, technologies and companies. Whether your goal is to research the latest datasheet from a company, speak directly with an Application Engineer, or jump to a company's technical page, the goal of **Get Connected** is to put you in touch with the right resource. Whichever level of service you require for whatever type of technology, **Get Connected** will help you connect with the companies and products you are searching for.

www.rtcmagazine.com/getconnected

Company	Page	Website
ACCES I/O Products	73	www.accesio.com
Advantech Technologies, Inc.	48	www.advantech.com
Ampro Computers, Inc.	35	www.ampro.com
Apacer Memory America, Inc.	15	www.apacer.com
ARCOM	34	www.arcom.com
BitMicro Networks, Inc.	6	www.bitmicro.com
BittWare,	46	www.bittware.com
Concurrent Technologies	67	www.gocct.com
Digital Logic AG	22	www.adlogic-pc104.com
Dynatem, Inc.	7	www.dynatem.com
ELMA Electronic, Inc.	32	www.elma.com
Embedded Planet	41	www.embeddedplanet.com
Enea Embedded Technology	19	www.ose.com
Fastwel Co Ltd	6	www.fastwel.com
GE Fanuc Embedded Systems	8	www.gefanuc.com/embedded
Hunt Engineering Ltd	68	www.hunt.rtg.com
Interactive Circuits and Systems	75	www.ics-ltd.com
Kontron America	13,76	www.kontron.com
KUKA Controls GmbH	42	www.kuka-controls.de
Mercury Computer Systems	58	www.mercury.com
Micro/sys, Inc.	29	www.embeddedsys.com
Microsoft Tour	54	www.microsoftembedded.com
Microsoft Windows Embedded	38,39	www.learnaboutembedded.com/shopping2
Octagon Systems	2,3	www.octagonsystems.com
One Stop Systems	60	www.onestopsystems.com
Phoenix International	68	www.phenxint.com
QNX Software Systems, Ltd.	11	www.qnx.com
Real-Time & Embedded Computing Conference	53,59	www.rtecc.com
Real-Time Innovations, Inc.	18	www.rti.com
Red Rock Technologies, Inc.	63	www.redrocktech.com
Ricoh	28	www.rei.ricoh.com
RTC Digital Delivery	70	www.rtcmagazine.com
SBE, Inc.	69	www.sbei.com
SBS Technologies	4	www.sbs.com
Swell Software	27	www.swellsoftware.com
TenAsys Corporation	24	www.tenasys.com
Themis Computer	49	www.themis.com
Ultimate Solutions, Inc.	23	www.ultsol.com
VadaTech	36	www.vadatech.com
VersaLogic Corporation	43	www.versalogic.com
WinSystems	30	www.winsystems.com

RTC (Issn#1092-1524) magazine is published monthly at 905 Calle Amanecer, Ste. 250, San Clemente, CA 92673. Periodical postage paid at San Clemente and at additional mailing offices. POSTMASTER: Send address changes to RTC, 905 Calle Amanecer, Ste. 250, San Clemente, CA 92673.

600 MBytes/second Record and Playback... ...Delivered!

In today's demanding data acquisition and sensor processing applications, speed is everything. More inputs, increasing channel densities and higher bandwidth place unprecedented demands on hardware. The new daqPC from ICS rises to this challenge with a leading 600 MBytes/second sustained, continuous record and playback speed — half as fast again as its predecessor, and way ahead of the competition.

But it's not just powerful. It is also flexible, compact and easy-to-use — which means that the range of sonar, radar and radio applications for which it can be used is limited only by your imagination.

ICS
SENSOR PROCESSING

Go ahead. Imagine.

CHECK THE FULL SPEC NOW AT
www.ics-ltd.com/daqPC

OR CALL TOLL FREE (US ONLY)
800-267-9794 or 613-749-9241

The design possibilities are **wide open**

open

OPEN MODULAR SOLUTIONS

ACCESS

EDGE

CORE

TRANSPORT

DATA CENTER

DESIGN AND DEPLOY

Your new IMS infrastructure applications using Kontron AdvancedTCA / AdvancedMC modular solutions.

Kontron simply takes the worry – and the expense – out of building complex IMS communication platforms for **fixed-mobile networks.** Whatever the application, you can design and deploy your project faster than you think with fully integrated, open standard modular solutions that are application-ready, right off the shelf.

That means reduced development costs for you, and tremendous **"swap in - swap out"** service flexibility for your carrier customers. It's a very smart win-win go-to-market strategy for everything from data and signaling platforms to IP streaming multimedia applications for video-on-demand, real-time voice and video telephony.

It's so simple. Start your next application with Kontron.

> Go Open Standards > Go Kontron > Ask for an Eval today >

www.kontron.com/ATCA

AT8001 Processor

AT8902 Base/Fabric Hub

AM4500

AM4001

AM4300

XL8000

AdvancedTCA®
AdvancedMC™

SERVICE AVAILABILITY FORUM

ASI SIG

Intel® Communications Alliance Premier Member

Computer On Modules	**Blades & Mezzanines**	CPU Boards	**Systems**	Mobile Rugged	Custom Solutions

kontron

Kontron and the Kontron logo are registered trademarks of Kontron AG. All other trademarks are the property of their respective owners. ©2006 Kontron America, Inc.

1-888-526-ATCA - sales@us.kontron.com — EMEA: +49 0 800 7253756 - ASIA: +886 2 2910 3532